Research on the Application and Ethical Boundaries of AI-Generated Content in Personalized Marketing Content Generation

Yilong Zhu*

Hainan Vocational University of Science and Technology, Haikou, 571126, China
**Corresponding author: 13637655987@163.com*

Abstract: *The rapid advancement of generative artificial intelligence is propelling personalized marketing into a new intelligent phase. This study systematically analyzes the technical mechanisms of AIGC-driven content generation, including semantic parsing based on large language models, visual generation through multimodal fusion, and the paradigm shift from static templates to dynamic interaction. It constructs a collaborative framework encompassing in-depth user profile development, brand semantic field calibration, and content ecosystem architecture, while conducting an in-depth examination of ethical boundaries concerning data privacy, algorithmic transparency, content authenticity, and cognitive autonomy. Through this dual-dimensional analysis of technology and ethics, the research provides theoretical foundation and practical guidance for responsible innovation in the application of AIGC within the marketing domain.*

Keywords: *generative artificial intelligence, personalized marketing, content generation, multimodal fusion, data privacy, algorithmic ethics*

Introduction

Against the backdrop of deepening digital transformation, personalized marketing has become a core strategy for enterprises to enhance user engagement and brand value. However, traditional content generation methods face significant bottlenecks in terms of efficiency, scale, and precision, struggling to meet user expectations for highly personalized content in dynamic market environments. The emergence of generative artificial intelligence technology provides a new technical pathway to break through this dilemma. Through its capabilities in semantic comprehension, multimodal generation, and dynamic optimization, it has achieved a paradigm shift in marketing content from standardized production to intelligent personalized creation.

The significance of this research lies in systematically deconstructing the technological implementation mechanisms and collaborative operational models of AI-generated content (AIGC) in personalized marketing, while simultaneously revealing the ethical challenges arising from its application processes. Through a dialectical analysis of technological empowerment and ethical constraints, this study attempts to construct a theoretical framework that balances innovative efficiency and social responsibility, thereby providing essential academic support for the scientific application and standardized development of AIGC in the marketing field.

1. Technological Mechanisms and Model Evolution of AIGC-Driven Personalized Content Generation

1.1 Semantic Comprehension and Dynamic Copy Generation Based on Large Language Models

The mechanism for semantic comprehension and dynamic copy generation based on large language models is rooted in the deep linguistic knowledge acquired through pre-training on large-scale corpora. Its technical core lies in the multi-layer self-attention network constructed by the Transformer architecture, which enables dynamic, non-local relational weight allocation between vocabulary and contextual compositions in the input text, thereby achieving deep semantic parsing that transcends superficial syntactic structures.

In personalized marketing scenarios, the system first encodes multi-dimensional user data,

including but not limited to historical search queries, social media interactions, browsing behavior sequences, and explicit feedback, forming a high-dimensional user intent vector. This vector serves as a guiding condition for the generation process and interacts with the universal language knowledge base inherent in large language models[1].

The model's generation process is essentially a sequential decision-making task based on probability distributions. During the decoding phase, strategies such as beam search and nucleus sampling are employed to select lexical sequences that best align with the current user context from the vast potential lexical space, while ensuring linguistic fluency and creativity.

This process not only achieves precise semantic mapping of users' personalized needs, but also enables style transfer and emotional adaptation. For instance, it can transform the functional description of the same product into either a technical exposition suitable for professionals or a relatable narrative intended for general consumers. Such end-to-end capability, spanning from semantic comprehension to content generation, signifies that marketing copy creation has transitioned from relying on manual rules and template matching to a data-driven, dynamically evolvable intelligent new stage.

1.2 Technical Pathway for Visual Content Generation under Multimodal Fusion

Visual content generation under multimodal fusion represents a significant leap in the perceptual and creative dimensions of AIGC. Its technical pathway is built upon a collaborative framework integrating cross-modal representation learning with generative adversarial networks or diffusion models. The core breakthrough lies in constructing a unified semantic space, enabling information from different modalities—such as textual descriptions, keyword tags, user preference vectors, and even sample images—to be mapped into shared, aligned mathematical representations.

Cross-modal pre-training techniques, represented by the CLIP model, utilize contrastive learning on massive image-text pairs to successfully embed feature representations of natural language instructions and visual concepts into the same vector space, thereby providing precise semantic navigation for text-to-image generation.

Building upon this foundation, technologies such as diffusion models accomplish image synthesis through a progressive denoising process. This process begins with a Gaussian-distributed random noise field, which is conditionally guided by the aforementioned cross-modal semantic vector. Through a series of iterative steps, noise is progressively removed, causing image structures, textures, and styles to emerge from chaos and ultimately converge into a clear visual representation that closely aligns with the textual instructions[2].

This pathway enables the system to comprehend complex and abstract personalized requirements such as "generate a product display image of sports shoes for young urban males combining minimalist aesthetics with a futuristic style, set against a neon-lit urban nightscape," and transform them into unique visual assets. It fundamentally transforms the traditional workflow reliant on stock image searches and manual design, thereby automating and personalizing the entire process from conceptualization to visual execution, while significantly expanding the boundaries and efficiency of creative expression.

1.3 The Paradigm Shift in Content Generation: From Static Templates to Dynamic Interaction

The paradigm shift driven by AIGC fundamentally represents an evolution from closed, pre-defined static content systems to open, dynamic interactive ecosystems based on data feedback loops. Traditional content personalization relies on pre-constructed limited content template libraries and simple rule engines, conducting A/B testing or finite combinations according to user tags, which inherently constrains its flexibility and adaptability. In contrast, the paradigm introduced by AIGC establishes a continuous learning cycle of "generate-evaluate-optimize."

In this dynamic system, the AIGC model serves as the content generator, initially producing outputs based on its offline-trained knowledge. After the generated content is deployed, its effectiveness is continuously monitored and quantitatively evaluated through a series of real-time embedded tracking data and user implicit feedback, such as click-through rates, conversion rates, dwell time, and interaction depth.

These performance metrics form a real-time data stream, which is transformed into reinforcement learning signals interpretable by the model. Through online learning or fine-tuning mechanisms, the

model dynamically adjusts its generation strategy, tending to replicate successful content patterns while avoiding ineffective or negatively received ones. This process transforms content generation strategy from a one-time static configuration into a dynamically evolving entity that synchronizes with the preference distribution of the user population. Consequently, the system acquires the capability to capture subtle shifts in socio-cultural trends at a macro scale and respond to individual users' interest drift at a micro scale, ultimately achieving co-adaptation and collaborative evolution between marketing content and the dynamic market environment.

2. Content Generation Collaborative Paradigm of AIGC in Personalized Marketing

2.1 In-depth User Profile Construction and Precise Content Matching Strategy

Traditional user profiling methodologies primarily rely on superficial correlations between static attribute labels and discrete behavioral records, exhibiting significant limitations in capturing users' dynamic intent evolution and underlying psychological motivations [3].

The deep construction paradigm driven by AIGC technology achieves panoramic integration and dynamic parsing of multi-source heterogeneous data by introducing advanced algorithms such as temporal modeling, graph convolutional networks, and deep representation learning. This architecture not only integrates users' explicit behavioral trajectories but, more importantly, infers users' latent demand structures, value orientation maps, and decision-making psychological models through technical means including semantic parsing, sentiment computation, and social network analysis of interactive contexts, content consumption preferences, and relational networks.

This deep construction process yields a dynamic user profile, which is essentially a continuously updated set of high-dimensional feature vectors. Each dimension within this set represents a specific latent trait of the user and its evolutionary pattern over time. This vector serves as a critical conditional input embedded into the generation process of the AIGC model. Leveraging architectures such as conditional generative adversarial networks or variational autoencoders, the model can synthesize personalized content that aligns closely with the target user's cognitive framework across multiple dimensions, including topic selection, narrative logic, emotional tone, and expressive style.

This fundamental shift from "describing user characteristics" to "simulating user cognition" marks the transition of content matching strategies from rule-based mechanical filtering to a new stage of data-driven intelligent emergence, achieving truly personalized content delivery for individual users.

2.2 Maintenance of the Brand Semantic Field and Consistency Calibration of Generated Content

In the context of large-scale personalized content generation, how to ensure that massive automated output maintains high consistency with the core brand value system, visual identity guidelines, and discourse expression style constitutes a fundamental challenge for AIGC applications.

The brand semantic field, as a complex system of meanings composed of core concepts, cultural symbols, stylistic elements, and narrative patterns, needs to be translated into a structured constraint system that is machine-recognizable and executable[4].

Maintaining the uniformity of this field requires the construction of a multi-layered calibration framework. At the technical level, this involves creating a brand knowledge graph to formally define the brand's core narrative elements, value propositions, stylistic boundaries, and expression taboos through ontological methods, thereby establishing a computable brand semantic space.

At the model level, a technical approach combining Parameter-Efficient Fine-Tuning (PEFT), prompt engineering, and reinforcement learning is adopted to inject brand semantic constraints as strong conditional signals into the generation process of the AIGC model. Consequently, the model must perform dynamic trade-offs and optimization between responding to user personalization needs and adhering to brand guidelines. The generated content must not only satisfy individual user preference characteristics but also strictly conform to the boundary conditions defined by the brand knowledge graph in terms of semantic meaning, emotional tendency, and aesthetic expression.

This refined calibration mechanism effectively resolves the potential issue of brand image fragmentation in large-scale personalized production. It ensures that each personalized communication, while fostering innovation, consistently reinforces the unity and coherence of brand equity, thereby providing technical assurance for maintaining stable brand identity in a dynamic market environment.

2.3 Systematic Architecture and Operation of Large-Scale Personalized Content Ecosystems

The construction and operation of large-scale personalized content ecosystems require a systematic architecture that deeply integrates data flows, algorithmic models, and business processes. This architecture essentially constitutes an intelligent closed-loop system equipped with perception, decision-making, generation, and optimization capabilities, whose core lies in achieving end-to-end intelligentization from content creation to distribution.

In terms of architectural design, the system typically comprises four collaboratively functioning layers: The perception layer employs distributed data collection technologies to acquire user interaction behaviors, environmental context, and market trend data in real time. The decision-making layer utilizes online learning techniques, such as multi-armed bandit algorithms and contextual bandit models, to dynamically formulate optimal content strategies and resource allocation plans. The generation layer consists of a series of specially optimized conditional AIGC models responsible for efficiently synthesizing multimodal content assets according to instructions. The feedback layer establishes a comprehensive content performance evaluation indicator system to monitor content dissemination effects and user engagement in real time, generating quantified evaluation signals for the generation strategy.

Operating this complex ecosystem hinges on establishing data-driven automated workflows and continuous optimization mechanisms. This includes implementing automated quality inspection and compliance screening for AIGC output, constructing online learning and model iteration pipelines based on real-time performance data, and developing intelligent regulation capabilities for overall system resource utilization and cost-effectiveness. This systematic architecture ultimately transforms content marketing from traditional manual creative operations into a self-optimizing, highly scalable intelligent organism that achieves qualitative leaps in operational efficiency while ensuring maximized marketing effectiveness.

3. Constructing Ethical Boundaries for AIGC Applications in Personalized Marketing

3.1 Data Privacy Empowerment and the Algorithmic Transparency Paradox

The data collection and processing mechanisms in personalized marketing scenarios have transcended the coverage scope of traditional data governance frameworks, exhibiting unprecedented complexity and opacity. AIGC systems construct highly inferential digital mirrors of individuals by integrating users' explicit behavioral data, implicit contextual information, and cross-platform digital footprints. This process frequently involves deep inference of users' latent intentions, emotional states, and even personality traits, with the depth and breadth of data mining significantly exceeding users' conventional privacy expectations and cognitive boundaries. The current data protection paradigm, centered on notice-and-consent, reveals structural limitations in practice when confronted with dynamically evolving and continuously expanding model training requirements. The growing asymmetry between user authorization mechanisms and the actual scope of data processing has catalyzed new demands for privacy protection.

The black-box nature of generative algorithms further exacerbates this ethical dilemma. The technical architecture based on deep neural networks constructs highly nonlinear complex mapping relationships between inputs and outputs, making its internal decision-making logic and data flow paths difficult to observe, interpret, and trace externally, thus creating a unique algorithmic transparency paradox: while the system can produce highly accurate and contextually appropriate personalized content, its generation basis and reasoning process lack necessary interpretability. This transparency deficit not only impedes users' understanding of how and to what extent their personal data is being used but also risks diminishing the effectiveness of traditional data auditing and regulatory mechanisms[5].

Against this backdrop, there is a need to explore innovative balancing solutions among technical interpretability research, privacy-enhancing technologies, and commercial confidentiality protection, thereby promoting the formation of a new data governance paradigm adapted to the characteristics of generative AI.

3.2 Authenticity, Bias, and Accountability of Generated Content

AI-generated content (AIGC) faces ontological challenges regarding its authenticity. Models produce highly realistic yet entirely fictional content based on statistical patterns, potentially blurring the boundary between objective facts and algorithmic constructions, which may consequently impact consumers' right to know and capacity for judgment. This trend of authenticity ambiguity is particularly pronounced in marketing contexts, encompassing both exaggerated depictions of product features and completely fabricated usage scenarios or effect demonstrations.

Simultaneously, inherent social biases present in the training data are systematically amplified through the model's data assimilation and pattern replication mechanisms, resulting in stereotypical representations and implicit discrimination based on demographic characteristics, cultural backgrounds, and socioeconomic status in the output content.

This embedding and amplification mechanism of bias possesses a considerably concealed nature. The model does not merely replicate biases present in the training data; rather, it extracts and reinforces existing social asymmetries through its internal representation learning process, thereby generating outputs that align with statistical patterns but may violate principles of fairness. Furthermore, when such generated content triggers legal disputes or negative societal effects, the issue of accountability demonstrates unprecedented complexity.

Traditional accountability frameworks, established on clear causal relationships and identifiable actors, are difficult to adapt to the distinctive characteristics of AIGC, namely the plurality of creative agents and algorithmic autonomy. This necessitates a re-examination of existing liability distribution principles, requiring the establishment of a reasonable shared responsibility mechanism among multiple stakeholders—including model developers, platform operators, content users, and regulatory agencies—while simultaneously exploring new regulatory paradigms and accountability principles tailored to the technical features of generative AI.

3.3 Concerns Regarding Mental Manipulation and the Defense of Autonomous Decision-Making Rights

While AIGC-driven hyper-personalized marketing enhances commercial conversion efficiency, it also raises profound ethical concerns regarding users' cognitive autonomy. By leveraging microscopic insights into users' psychological characteristics, behavioral patterns, and decision-making preferences, the system can generate highly psychologically targeted persuasive content. This precision influence may transcend the boundaries of traditional marketing persuasion and approach ethical boundaries concerning cognitive freedom[6].

When personalized content strategies deliberately exploit users' emotional vulnerabilities, cognitive shortcuts, or subconscious motivations, their persuasive processes may bypass users' rational deliberation mechanisms, posing substantial challenges to consumers' autonomous decision-making rights.

Such implicit cognitive influence requires rigorous scrutiny through the lens of behavioral ethics. Research in neuroscience and behavioral economics indicates that precision persuasion based on personalized data may activate automated decision-making pathways in the human brain, thereby reducing the engagement of critical thinking.

In commercial contexts, the application of such technology without necessary ethical constraints may result in consumers making decisions that do not align with their long-term interests without being fully informed. Therefore, it is imperative to establish a dynamic balancing mechanism between technological innovation and ethical constraints, including developing impact assessment frameworks, setting transparent content disclosure standards, and exploring technical implementation pathways that respect users' cognitive autonomy. This requires technology developers and commercial practitioners not only to focus on maximizing algorithmic efficiency but also to consider its social implications and ethical boundaries, ensuring that technological advancement does not come at the expense of users' fundamental cognitive rights.

Conclusion

This study, through a three-dimensional analysis of the technical pathways, collaborative paradigms,

and ethical boundaries of AIGC in personalized marketing content generation, reveals its transformative potential from instrumental assistance to systemic reconstruction. Technically, semantic parsing based on large language models, multimodal generation, and dynamic interaction mechanisms collectively propel content generation toward intelligent and self-adaptive development. At the collaborative level, in-depth user profiling, brand semantic constraints, and systematic operational architecture form the foundation for achieving scalable personalization. Ethically, issues such as data privacy, algorithmic transparency, content authenticity, and cognitive autonomy urgently require the establishment of a multidimensional governance framework. Future research could further explore the integration pathways of AIGC with privacy-enhancing technologies, develop generative models with high interpretability and fairness, and construct an interdisciplinary collaborative ethical governance system to promote a dynamic balance between technological innovation and ethical constraints in the application of AIGC within the field of personalized marketing.

References

[1] Lin Xue. "Application of Generative AI in Personalized Marketing." Marketing World 14 (2025): 55-57.
[2] Huang Xuan. "AI Empowers the 'E-commerce Visual Marketing Design' Course in Vocational Colleges: Reshaping Teaching Content and Practice Models." China Electronic Commerce 31.11 (2025): 25-27.
[3] Jiao Pengfei. "The Impact of Generative AI-Based Personalized Marketing Content on Consumer Decision-Making." Marketing World 11 (2025): 58-60.
[4] Wang Shaoyi, et al. "Application of Large AI Models in Content Generation for Marketing Scenarios." Digital Technology and Applications 43.01 (2025): 102-104.
[5] Li Xueqin. "The Influence Mechanism and Countermeasures of AI-Generated Marketing Content on Consumer Approach/Avoidance Behaviors." Social Science Journal of Harbin Normal University 16.03 (2025): 104-108.
[6] Duan Chunlin and Chen Ruijie. "Generative AI Content Creativity and Marketing Value Evolution in the Fully Intelligent Era." Modern Advertising 04 (2024): 4-9.