IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 5, April 2026
Impact Factor: 8.2

The Use of Artificial Intelligence for Personalized Advertising and Marketing.

Dhruv Jivrajani, Shrihari Moharir, Dr. Surekha Gaikwad
MAEER MIT ACSC, Pune
dhruvjivrajani7@gmail.com, shriharimoharir7@gmail.com
slgaikwad@mitacsc.ac.in

Abstract: *Artificial Intelligence (AI) has fundamentally and irreversibly transformed the advertising and marketing landscape by enabling hyper-personalised consumer experiences at unprecedented scale and speed. This paper provides a comprehensive examination of how AI technologies - including machine learning, deep learning, natural language processing, computer vision, reinforcement learning, and generative AI - are being deployed to tailor marketing messages, product recommendations, pricing strategies, and creative content to individual consumers. It traces the historical evolution of personalised marketing from rudimentary rule-based systems to today's sophisticated neural architectures. It analyses the technical mechanisms underlying modern personalisation pipelines, surveys applications across major marketing channels, quantifies documented benefits, and critically examines the ethical, social, and regulatory challenges that AI-driven advertising presents. Finally, it charts the emerging trends and technologies that will define the next generation of personalised marketing, and proposes a framework for responsible deployment. The paper concludes that while AI offers transformative commercial opportunity, its long-term legitimacy depends on governance structures that centre consumer trust, transparency, and equity.*

Keywords: *Artificial Intelligence*

I. INTRODUCTION

Marketing has always been fundamentally about communication - the attempt by producers to reach potential consumers with messages that resonate, persuade, and ultimately drive behaviour. For most of commercial history, this was a blunt instrument. Billboards reached everyone who passed. Television advertisements interrupted whoever was watching. Even early digital advertising, for all its apparent sophistication, was largely a matter of placing banners on websites that attracted the right general demographic.

The advent of data-driven marketing, and especially AI-powered personalisation, has changed this calculus entirely. Today, marketing systems do not merely target demographic groups - they target individuals, in specific moments, on specific devices, with messages calibrated to their unique history, preferences, psychological profile, and predicted future behaviour. The aspiration, increasingly realised, is the delivery of the right message, to the right person, at the right time, through the right channel - automatically, at scale, and in real time.

The scale of this transformation is reflected in the numbers. Global digital advertising expenditure surpassed $600 billion in 2023 and is projected to exceed $870 billion by 2027, with AI-powered systems accounting for a growing majority of that spend. AI personalisation has become the central engine of the digital economy, underpinning the business models of the world's most valuable companies - Google, Meta, Amazon, Alibaba, and others - and increasingly adopted by businesses of every size and sector.

Yet the same capabilities that make AI marketing so commercially powerful also make it potentially intrusive, manipulative, and harmful. The deployment of AI to understand and influence individual human behaviour at scale raises profound questions about privacy, autonomy, fairness, and the appropriate boundaries of commercial persuasion.

DOI: 10.48175/IJARSCT-32854

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

These questions are not merely academic - they are being actively contested in courts, legislatures, and public discourse around the world.

This paper aims to provide a comprehensive, balanced, and technically grounded analysis of AI in personalised advertising and marketing. It is structured to serve both practitioners seeking to understand the technology and its applications, and researchers and policymakers seeking to understand its implications and governance challenges.

II. HISTORICAL EVOLUTION OF PERSONALISED MARKETING

2.1 Pre-Digital Era: Mass Marketing and Segmentation

The dominant paradigm of twentieth-century marketing was mass communication. Brands such as Coca-Cola, Ford, and Procter & Gamble built enormous enterprises by broadcasting consistent messages to the broadest possible audiences through print, radio, and television. The underlying assumption was one of homogeneity - consumers within a broad demographic were treated as interchangeable.

The concept of market segmentation, popularised by Wendell Smith in the 1950s, introduced the idea that consumers could be divided into meaningful subgroups based on shared characteristics, and that different groups might respond to different messages. This was a significant conceptual advance, but practical implementation remained limited by the cost of data collection and the constraints of broadcast media.

Direct marketing, pioneered in the mid-twentieth century, represented the first serious attempt at individualised communication. Catalogue retailers, insurance companies, and magazine publishers built mailing lists and began experimenting with personalised letters and targeted offers. Response rate analysis - the forerunner of modern A/B testing - was used to optimise these campaigns. However, the scale of individualisation remained modest, constrained by manual processes and limited data.

2.2 Early Digital Marketing: The First Wave of Data

The commercialisation of the internet in the 1990s transformed the possibilities for personalised marketing. For the first time, it was possible to observe individual consumer behaviour at scale - which pages they visited, which links they clicked, how long they spent on content, and what they purchased. This behavioural data was richer and more actionable than anything available to traditional marketers.

Amazon's launch of its collaborative filtering recommendation engine in 1998 - the "customers who bought X also bought Y" system - was a landmark moment. It demonstrated that algorithmic analysis of purchase data could surface genuinely relevant product suggestions, and that consumers would respond positively to them. Amazon has credited personalised recommendations with driving 35% of its total revenue at peak periods.

Early banner advertising on the web was largely contextual - ads for camera equipment appeared on photography websites, ads for financial products on news sites. Click-through rates were initially high, reflecting novelty, but declined rapidly as consumers became habituated to online advertising. This created pressure to improve targeting precision.

The introduction of browser cookies enabled websites to recognise returning visitors and track behaviour across sessions. Third-party cookies extended this tracking across different websites, enabling ad networks to build profiles of individual users' browsing behaviour even when they visited unrelated sites. This was the foundation of behavioural targeting - a major step toward individualised advertising, but still primitive by contemporary standards.

2.3 The Programmatic Revolution

The mid-2000s and early 2010s saw the rise of programmatic advertising - the automated buying and selling of digital ad impressions through real-time bidding (RTB) systems. In a programmatic auction, when a user loads a webpage, an auction is conducted in the milliseconds before the page renders. Advertisers submit bids based on the value they assign to reaching that particular user, and the highest bidder's ad is displayed.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

Programmatic advertising vastly increased the efficiency of digital ad markets, enabling advertisers to target specific audience segments across millions of websites simultaneously. Demand-Side Platforms (DSPs) allowed advertisers to manage bids across multiple exchanges; Supply-Side Platforms (SSPs) allowed publishers to manage their inventory. Data Management Platforms (DMPs) aggregated audience data from multiple sources to enrich user profiles.

However, early programmatic systems were still relatively unsophisticated in their use of data. Targeting was largely based on keyword associations, demographic inferences, and broad interest categories. The models were not deeply personalised - they targeted segments, not individuals.

2.4 The Deep Learning Era

The breakthrough achievements of deep learning in the early 2010s - beginning with AlexNet's dramatic improvement in image classification in 2012 - fundamentally altered what was computationally achievable in AI. Researchers demonstrated that deep neural networks could learn rich, abstract representations from raw data, outperforming hand-engineered feature extraction across a wide range of domains.

These advances rapidly propagated into commercial AI systems, including advertising and marketing. Deep learning enabled far more sophisticated models of consumer behaviour, capable of processing higher-dimensional data, identifying non-obvious patterns, and making more accurate predictions of consumer intent and likely response.

Facebook (now Meta) was among the first to deploy deep learning at scale for ad targeting. Its early deep learning-based news feed ranking system, introduced around 2015, used neural networks to personalise the content each user saw. Google similarly integrated deep learning into its Search and Display ad ranking systems. The result was a significant improvement in ad relevance and engagement metrics across both platforms.

The introduction of transformer architectures (Vaswani et al., 2017) and the subsequent development of large language models (LLMs) - BERT, GPT series, and others - opened new frontiers. These models enabled far more nuanced understanding of natural language, allowing systems to interpret search queries, product reviews, social media content, and customer service interactions with a depth previously unattainable.

2.5 The Present: AI-Native Marketing

Today, AI is not a tool applied to marketing - it is the infrastructure of digital marketing. Every major platform's advertising system is fundamentally an AI system. Every impression served, every bid placed, every creative selected, and every audience modelled involves machine learning at its core. Increasingly, the creative content of ads itself is being generated by AI.

This represents a qualitative shift. Earlier generations of personalised marketing used data to select among a fixed set of creative options or to identify audience segments. Contemporary AI marketing systems can generate unique messages, offers, and creative assets for individual users, and can continuously adapt in response to each interaction.

III. CORE AI TECHNOLOGIES IN PERSONALISED MARKETING

3.1 Supervised Machine Learning

Supervised learning - training models on labelled historical data to predict outcomes - is the workhorse of marketing AI. Applications include:

Propensity Modelling: Predicting the probability that a given user will take a specific action (click an ad, make a purchase, subscribe to a service, or churn). These models enable prioritisation of marketing effort toward high-probability targets.

Lead Scoring: Assigning scores to potential customers based on predicted likelihood to convert, enabling sales teams to prioritise outreach.

Lookalike Modelling: Identifying users whose characteristics resemble those of existing customers, expanding addressable audiences while maintaining relevance. Meta and Google both offer lookalike audience tools based on this approach.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

Churn Prediction: Identifying customers at risk of discontinuing a relationship, enabling proactive retention campaigns.

Common algorithms include logistic regression, gradient boosted trees (XGBoost, LightGBM), and deep neural networks, depending on the complexity of the prediction task and the volume of data available.

3.2 Deep Learning and Neural Networks

Deep learning enables the processing of high-dimensional, unstructured data - images, video, audio, text - that is not amenable to traditional ML methods. In marketing contexts:

Embedding Models: Represent users and items (products, content, ads) as dense vectors in a shared high-dimensional space. Items and users with similar characteristics are positioned close together, enabling efficient similarity search for recommendations and targeting.

Sequential Models (RNNs, LSTMs, Transformers): Model the temporal sequence of user interactions, capturing how preferences and intent evolve over a session or over time. Particularly valuable for modelling browsing journeys and predicting purchase intent.

Deep Neural Networks for CTR Prediction: Click-through rate prediction - estimating the probability a given user will click a given ad - is one of the most commercially important prediction tasks in digital advertising. Deep learning models significantly outperform earlier approaches on this task.

3.3 Natural Language Processing

NLP enables machines to understand and generate human language, with wide-ranging marketing applications:

Sentiment Analysis: Classifying the emotional tone of consumer-generated text - reviews, social posts, customer service interactions - as positive, negative, or neutral. Enables brands to monitor consumer perception and identify issues in real time.

Intent Detection: Identifying what a consumer is seeking from their search query or chatbot input, enabling more relevant ad serving or content recommendation.

Topic Modelling: Discovering latent themes in large corpora of text (reviews, social media, support tickets), surfacing consumer concerns and interests at scale.

Named Entity Recognition: Extracting structured information - brand names, product categories, locations - from unstructured text.

Dynamic Content Generation: LLMs can generate personalised ad copy, email subject lines, and product descriptions tailored to individual user profiles, dramatically expanding the scale at which creative personalisation is possible.

3.4 Computer Vision

Computer vision enables AI systems to interpret and generate visual content:

Visual Preference Modelling: Analysing the images and videos a user engages with to infer aesthetic preferences and lifestyle indicators. Pinterest and Instagram use this to power visually matched product recommendations.

Product Recognition: Identifying products within user-generated content, enabling shoppable posts and targeted ads based on visual interest signals.

Ad Creative Analysis: Automatically tagging and categorising creative assets (images, videos) by their visual attributes, enabling intelligent matching of creative to audience.

Generative Visual AI: Diffusion models and GANs can generate unique visual creative assets for ads, enabling image-level personalisation.

3.5 Recommender Systems

Recommender systems are the most commercially visible application of AI personalisation:

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

Collaborative Filtering: Identifies patterns of similarity across users' interaction histories. Users who have similar past behaviour are likely to have similar future preferences. Does not require knowledge of item content - purely data-driven.

Content-Based Filtering: Recommends items similar in content (category, attributes, description) to items the user has engaged with previously. Less dependent on having a large user base.

Hybrid Systems: Combine both approaches to overcome limitations of each. Most mature recommender systems (Netflix, Spotify, Amazon) use sophisticated hybrid architectures.

Session-Based Recommendation: Models short-term intent based on actions within a current session, complementing longer-term preference models.

Knowledge Graph Enriched Recommendations: Incorporate structured knowledge about relationships between items, brands, and concepts to improve recommendation quality, especially for new items or users with limited history.

3.6 Reinforcement Learning

Reinforcement learning (RL) treats the ad delivery problem as a sequential decision-making task, where the system learns a policy that maximises long-term reward (typically conversions or revenue) through interaction with users:

Contextual Bandits: A simplified RL formulation where the system selects which ad to show (the "arm") based on contextual features of the user and situation, and learns from the observed reward (click or no click). Widely used in real-time ad serving.

Deep RL for Bidding: Systems that learn optimal bidding strategies in programmatic auctions by treating each auction as a step in a long-term optimisation problem. Can account for budget constraints, cross-channel effects, and customer lifetime value.

Exploration vs. Exploitation: RL systems must balance exploiting known-good strategies with exploring new ones that might perform better. This is a fundamental challenge in commercial deployment, where the cost of suboptimal decisions is real.

3.7 Generative AI

The most recent and rapidly developing category of AI marketing technology:

Large Language Models for Copy Generation: GPT-4, Claude, Gemini, and similar models are being used to generate ad headlines, body copy, email content, product descriptions, and social media posts at scale. Systems like Persado and Phrasee use LLMs fine-tuned on marketing data to generate and test copy variants.

Image and Video Generation: Diffusion models (DALL-E, Stable Diffusion, Midjourney, Sora) enable the generation of visual ad creative without requiring human designers or photographers. Brands can generate thousands of creative variants for testing, or create unique visuals for individual users.

Personalised Video: Emerging systems can generate or modify video content to incorporate individual-specific elements - a customer's name, local imagery, or personalised offers - at scale.

IV. THE AI PERSONALISATION PIPELINE

Understanding AI-driven personalised marketing requires understanding the end-to-end technical pipeline that transforms raw data into delivered ad experiences.

4.1 Data Collection and Sources

The foundation of personalised marketing is data. Contemporary systems draw on an increasingly diverse range of data sources:

First-Party Data: Collected directly by the brand through its own owned channels. Includes purchase history, account information, website behaviour, app usage, email engagement, and customer service interactions. First-party data is the highest-quality and most ethically straightforward source.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 5, April 2026
Impact Factor: 8.2

Second-Party Data: Another company's first-party data, shared through a direct partnership. For example, an airline might share travel data with a hotel chain under a co-marketing agreement.

Third-Party Data: Aggregated by data brokers from a wide range of sources, including public records, loyalty programmes, financial data, and purchased datasets. Third-party data dramatically expands the profile that can be built on a consumer, but raises significant privacy concerns and is increasingly restricted by regulation and browser policy.

Behavioural Data: Real-time tracking of digital behaviour - page views, clicks, searches, video watches, scroll depth, time on page. Provides a dynamic, intent-rich signal about current consumer interests.

Social Data: Publicly accessible social media activity, engagement patterns, and network connections. Enables inference of interests, life events, and social influence.

Contextual Data: Information about the environment in which an ad is being served - the content of the page, the device type, time of day, location. Used to ensure ad relevance to the immediate context.

Offline Data: Purchase transactions, in-store visits, call centre records. Increasingly linked to digital profiles through identity resolution.

4.2 Identity Resolution and Data Integration

One of the most technically challenging aspects of personalised marketing is assembling data from disparate sources into a coherent individual profile. A single consumer may be known by different identifiers across different systems - an email address in a CRM, a device ID in a mobile app, a cookie ID on a website, a loyalty number in a retail system.

Customer Data Platforms (CDPs) are purpose-built systems for unifying these fragmented identity signals into a single customer profile. They use deterministic matching (identical email addresses or phone numbers) and probabilistic matching (inferred from behavioural patterns, device characteristics, and timing) to link records across systems.

The deprecation of third-party cookies - announced by Google for Chrome and already implemented in Safari and Firefox - is significantly disrupting this process, forcing the industry toward first-party identity solutions and privacy-preserving alternatives.

4.3 Feature Engineering and Representation Learning

Raw data must be transformed into representations suitable for machine learning models:

Feature Engineering: Domain-knowledge-driven construction of input variables from raw data. For example, deriving "days since last purchase," "product category affinity score," or "price sensitivity index" from transaction records.

Embedding Learning: Neural networks can learn dense vector representations of users, products, and contexts directly from interaction data, capturing complex latent relationships without manual feature engineering.

Real-Time Feature Computation: For real-time personalisation, features must be computed in milliseconds at the time of each ad impression. This requires sophisticated streaming data infrastructure.

4.4 Model Training and Updating

Models are trained on historical data to learn associations between features and outcomes. Key considerations include:

Training Data Volume: Deep learning models require large training datasets. Large platforms have significant advantages here - Meta processes trillions of data points; smaller advertisers must rely on platform-provided models.

Label Definition: The choice of training signal (clicks, conversions, revenue, long-term retention) has profound effects on model behaviour and user experience. Optimising for clicks may not optimise for genuine consumer satisfaction.

Model Freshness: Consumer behaviour is non-stationary - trends, interests, and contexts change. Models must be retrained or updated regularly to remain effective.

Feedback Loops: When models influence which ads are shown, they influence which interactions are observed, which influence future model training. These feedback loops can cause models to become self-reinforcing and increasingly narrow.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 5, April 2026
Impact Factor: 8.2

4.5 Real-Time Inference and Serving

At the moment a user encounters an ad placement - whether on a website, in an app, or in search results - the personalisation system must:

Retrieve the user's current profile and features

Generate candidates (ads potentially relevant to this user)

Score candidates using one or more ranking models

Apply business rules (budget caps, frequency limits, brand safety filters)

Return the selected ad for display

This entire process occurs in 50–200 milliseconds in programmatic advertising - the time available before the page renders. At platform scale (Google processes over 8.5 billion searches per day; Meta serves over 3 billion users), this represents an extraordinary engineering challenge.

4.6 Measurement and Attribution

Determining the causal effect of advertising on consumer behaviour is one of marketing's most enduring challenges. AI has improved attribution, but fundamental difficulties remain:

Multi-Touch Attribution: ML models can assign fractional credit for a conversion to each marketing touchpoint in a consumer's journey, replacing simplistic last-click attribution with more nuanced causal models.

Incrementality Testing: Randomised controlled experiments (holdout tests) measure the true incremental lift attributable to advertising by comparing behaviour between exposed and unexposed groups.

Media Mix Modelling: Statistical models that decompose aggregate sales into contributions from different marketing channels, incorporating data from both online and offline sources.

V. APPLICATIONS ACROSS MARKETING CHANNELS AND INDUSTRIES

5.1 Search Advertising

Search advertising remains the largest segment of digital advertising by revenue. Google's AI-powered search advertising system combines several layers of personalisation:

Query Understanding: NLP models interpret the meaning and intent behind search queries, going beyond keyword matching to understand conceptual relevance.

Smart Bidding: Google's automated bidding system uses ML to set bids in real time for each auction based on the predicted probability of conversion, incorporating signals including device, location, time of day, search history, and many other contextual factors. Advertisers set target outcomes (ROAS, CPA); the AI manages execution.

Responsive Search Ads: Advertisers provide multiple headline and description options; Google's AI learns which combinations perform best for different users and contexts, assembling personalised ad text dynamically.

Performance Max: An AI-driven campaign type that automatically distributes budget and creative across all Google surfaces (Search, Display, YouTube, Gmail, Maps) to maximise conversions.

5.2 Social Media Advertising

Social platforms have built some of the world's most sophisticated AI personalisation systems, using the rich behavioural and social graph data they accumulate:

Meta (Facebook and Instagram): Meta's ad system models hundreds of signals per user, including on-platform behaviour, off-platform browsing (via the Meta Pixel), social connections, and declared interests. Its Advantage+ Shopping Campaigns use AI to automate audience targeting, bidding, and creative selection. Meta's AI systems have faced significant regulatory scrutiny, particularly in the EU, for their use of off-platform data.

TikTok: TikTok's For You Page algorithm is widely regarded as one of the most effective content personalisation systems ever built. Its ad platform leverages the same infrastructure, using content engagement signals to target ads with remarkable precision, particularly effective at reaching younger demographics.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

LinkedIn: Specialises in professional context-aware targeting - job title, company, industry, career stage - enabling B2B marketers to reach decision-makers with unusual precision.
Pinterest: Visual discovery platform where AI matches products to users based on the aesthetic and thematic patterns of their saved content. Shoppable Pins enable direct commerce from personalised recommendations.

5.3 Programmatic Display and Video Advertising

The programmatic ecosystem - spanning display banners, native ads, connected TV, audio, and digital out-of-home - is almost entirely AI-driven:
Audience Targeting: DSPs use ML models to identify and bid on users matching specified audience profiles across millions of websites and apps.
Contextual Targeting: With cookie deprecation, AI-powered contextual targeting - understanding the semantic content of a page to serve relevant ads - is experiencing a renaissance. Systems go far beyond keyword matching to understand page meaning, tone, and suitability.
Brand Safety: ML classifiers analyse page content in real time to prevent ads from appearing alongside inappropriate content.
Programmatic Creative Optimisation (PCO): AI selects and assembles ad creative components (headline, image, call-to-action) dynamically to create personalised combinations for each impression.

5.4 Email Marketing

Email remains one of the highest-ROI marketing channels, and AI has substantially enhanced its effectiveness:
Send Time Optimisation: ML models predict the time each individual subscriber is most likely to open and engage with email, scheduling delivery accordingly.
Subject Line Optimisation: AI generates and tests multiple subject line variants, learning which linguistic patterns drive highest open rates for different audience segments.
Personalised Content Blocks: Email templates contain dynamic sections populated with AI-selected content - product recommendations, articles, offers - personalised to each recipient.
Automated Journey Optimisation: AI manages the sequencing, timing, and content of email communications across complex multi-step customer journeys, adapting based on each individual's responses.

5.5 E-Commerce and Retail

E-commerce is perhaps the domain where AI personalisation has achieved its most visible and measurable impact:
Homepage and Category Personalisation: Product displays are ordered and filtered based on individual user preferences, presenting the most relevant items first.
Dynamic Pricing: AI models set prices that vary by user, time, competitive context, and inventory levels. While controversial from an equity standpoint, dynamic pricing is now standard in sectors including travel, ride-hailing, and live events.
Search Relevance: On-site product search is personalised to each user's purchase history and preferences, going beyond keyword matching to understand intent.
Cart Abandonment and Retargeting: ML models identify users who have abandoned shopping carts and trigger personalised re-engagement campaigns with tailored messaging and incentives.
Post-Purchase Experience: Personalised follow-up communications, loyalty offers, and cross-sell recommendations extend the customer relationship beyond the initial transaction.

5.6 Streaming and Entertainment

Content personalisation in streaming represents AI at its most sophisticated and commercially consequential:

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

Netflix: Uses multiple ML models to personalise every aspect of the user experience - what content is recommended, in what order, and even which thumbnail image is displayed for each title (with different thumbnails shown to different users to maximise appeal). Netflix has estimated that its recommendation system saves approximately $1 billion annually in reduced churn.

Spotify: Combines collaborative filtering, content-based audio analysis, and NLP (analysis of music press and fan blogs) to generate personalised playlists (Discover Weekly, Daily Mixes) that are widely credited with driving high user engagement and retention.

YouTube: Balances watch time optimisation with creator monetisation and advertiser brand safety through a multi-objective ML ranking system processing signals from billions of daily interactions.

5.7 Financial Services

AI personalisation in financial services encompasses insurance pricing, credit offers, investment recommendations, and fraud detection:

Personalised Insurance Pricing: Telematics data (driving behaviour), health and lifestyle data, and AI models enable highly individualised risk assessment and pricing, far more granular than traditional actuarial methods.

Next Best Offer: Banks use ML to predict which financial products (mortgage, loan, credit card, investment account) each customer is most likely to need and respond to at a given moment.

Personalised Financial Planning: Robo-advisors use AI to construct and manage individualised investment portfolios based on customer goals, risk tolerance, and time horizon.

VI. DOCUMENTED BENEFITS OF AI-DRIVEN PERSONALISED MARKETING

6.1 Commercial Performance

The business case for AI personalisation is well-supported by empirical evidence:

Research from McKinsey & Company has found that personalisation can reduce customer acquisition costs by up to 50%, lift revenues by 5–15%, and increase marketing ROI by 10–30%. Their 2021 research found that companies that excel at personalisation generate 40% more revenue from those activities than average players.

Amazon's recommendation engine has been credited with generating 35% of its total revenue. Netflix attributes its low churn rate in part to the effectiveness of its recommendation system, which it values at over $1 billion annually in retained subscription revenue. Spotify's Discover Weekly playlist has been cited as one of the most significant drivers of user engagement, with hundreds of millions of streams per week.

Segmentation studies consistently show personalised email campaigns achieving open rates 26% higher and click-through rates 14% higher than non-personalised equivalents. Personalised website experiences have been shown to increase conversion rates by 10–20% in controlled experiments.

6.2 Consumer Experience Improvement

When executed well, personalisation genuinely improves consumer experience. Consumers are presented with products, content, and offers relevant to their interests, reducing the cognitive load of navigating large product catalogues or content libraries. The discovery function of recommendation systems - helping consumers find things they would enjoy but would not have found independently - is a legitimate value creation activity.

Survey research consistently shows that consumers prefer relevant advertising to generic advertising, and that they are more likely to feel positively toward brands that demonstrate understanding of their needs. The Salesforce "State of the Connected Customer" report has found that the majority of consumers expect brands to personalise interactions.

6.3 Operational Efficiency

AI-driven marketing automation delivers substantial efficiency gains for marketing organisations:

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

Campaign management that would previously have required large teams of analysts and creative professionals can now be largely automated, with AI handling audience selection, bid management, creative testing, and reporting. This enables smaller teams to manage more complex, multi-channel campaigns at greater scale.

AI-powered attribution and measurement provide clearer insight into which marketing activities are driving commercial outcomes, enabling more intelligent budget allocation and reducing waste.

Creative production costs are reduced as generative AI enables the production of large numbers of creative variants for testing without proportional increases in production cost.

VII. ETHICAL CHALLENGES AND SOCIETAL CONCERNS

7.1 Privacy and Surveillance

The data requirements of AI personalisation necessitate pervasive surveillance of consumer behaviour. The volume, granularity, and sensitivity of data collected by major advertising platforms is extraordinary - every search query, every page visited, every post liked, every video watched, every physical location visited, every purchase made.

The aggregation problem - the ability to combine individually innocuous data points into deeply revealing profiles - is particularly concerning. From browsing history alone, it is possible to infer health conditions, financial status, political views, sexual orientation, relationship status, and psychological vulnerabilities. The Cambridge Analytica scandal demonstrated the potential for misuse of such profiles at political scale.

The transition away from third-party cookies, while positive for privacy, has not reduced the volume of data collection - it has shifted it toward first-party data (collected directly by platforms) and privacy-invading alternatives such as browser fingerprinting and email hashing, which are harder for consumers to detect or control.

7.2 Psychological Manipulation and Exploitation

The explicit goal of AI personalisation in advertising is to increase the probability of consumer action - clicks, purchases, sign-ups, donations. When this involves matching message to moment (showing ads for food delivery to hungry users late at night, or showing alcohol ads to users who have recently posted about stress), it crosses from relevance into exploitation of psychological vulnerabilities.

Research in behavioural economics has documented the many ways in which human decision-making deviates from rationality - anchoring, loss aversion, social proof, scarcity effects, and many others. AI systems can identify and exploit these biases at the individual level, deploying different persuasion strategies for different psychological profiles. While this is commercially effective, its compatibility with consumer autonomy and informed consent is questionable.

The use of AI in advertising to children is a particular concern. Children's cognitive development limits their ability to recognise and resist commercial persuasion. Many jurisdictions have legal protections against advertising to children, but AI systems can identify child users through behavioural signals and serve them age-appropriate content in ways that may circumvent formal restrictions.

7.3 Filter Bubbles and Reduced Serendipity

Recommendation and personalisation algorithms, by definition, show users more of what they have already engaged with. This creates the risk of filter bubbles - progressively narrowing information environments in which consumers are only exposed to products, ideas, and perspectives that reinforce their existing preferences and biases.

In commerce, this can manifest as brand lock-in - AI systems that consistently recommend familiar brands reduce the likelihood of consumers discovering alternatives, potentially reducing competition and consumer welfare. In content recommendation, there are concerns about algorithmic amplification of extreme content, where engagement-optimised recommendation systems have been observed to push users toward progressively more radical content.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

7.4 Algorithmic Bias and Discrimination

AI models trained on historical data inherit and can amplify the biases present in that data. In advertising contexts, this has led to documented cases of discriminatory ad delivery:

Research by Pro Publica and academic researchers has found that housing and employment ads on Facebook were delivered in racially skewed patterns - white users were more likely to see ads for certain housing options; women were underrepresented in ads for high-paying jobs in STEM fields. The platform's own optimisation systems, seeking to maximise engagement, replicated and amplified existing societal patterns of discrimination.

These outcomes may occur without any intent to discriminate on the part of advertisers or platform operators - they emerge from the interaction of biased training data, engagement-optimised objectives, and the self-reinforcing dynamics of algorithmic systems.

7.5 Transparency and Explainability

Most high-performing AI personalisation systems are effectively black boxes - their internal decision-making processes are not interpretable even by their designers. A consumer who sees an unexpected or unwanted ad has no meaningful ability to understand why they were targeted, and often no effective recourse.

This opacity creates multiple problems. Consumers cannot make informed decisions about data sharing when they do not understand how their data will be used. Advertisers cannot fully explain why their campaigns reached certain audiences. Regulators struggle to enforce compliance with non-discrimination requirements when algorithmic processes are opaque. And internal oversight within companies is hampered when engineers themselves cannot explain model behaviour in individual cases.

7.6 Concentration of Power

The infrastructure of AI-powered advertising is highly concentrated. Google and Meta together account for approximately 50% of global digital advertising revenue. This concentration reflects genuine network effects and data advantages - these platforms have vast first-party data that smaller competitors cannot match - but it also raises concerns about market power, barriers to entry, and the bargaining power of advertisers and publishers.

The walled garden model - where platforms collect user data, build models, and run ads, but share minimal data with advertisers - means that brand advertisers are largely dependent on platform-provided measurement and verification, creating conflicts of interest that have in several cases led to metric inflation scandals.

VIII. REGULATORY AND LEGAL LANDSCAPE

8.1 General Data Protection Regulation (GDPR)

The EU's General Data Protection Regulation, effective from May 2018, established the most comprehensive data protection framework in the world. Key provisions relevant to personalised advertising include:

The requirement for a lawful basis for processing personal data - most commonly, explicit consent or legitimate interests. Consent must be freely given, specific, informed, and unambiguous. The widespread practice of consent via pre-ticked boxes or buried settings was explicitly prohibited.

The right to object to processing for direct marketing purposes, without requirement to provide justification. The right to explanation of automated decisions that significantly affect an individual. Restrictions on processing of special categories of data (health, sexual orientation, political views, religion) - precisely the categories that AI profiling tends to infer.

GDPR enforcement has been significant. Meta has been fined over €4 billion in GDPR penalties as of 2024 in multiple rulings that found its advertising data practices unlawful. Google has faced fines in France, Spain, and Italy. The regulation has materially changed data practices across the industry.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

8.2 California Consumer Privacy Act (CCPA) and CPRA

The CCPA, effective from January 2020 and strengthened by the CPRA amendments effective 2023, established US consumers' rights regarding their personal data: the right to know what data is collected and how it is used; the right to delete personal data; the right to opt out of the sale or sharing of personal data; and the right to non-discrimination for exercising these rights.

California's law has had national impact, as many companies have implemented CCPA-compliant practices across their entire US customer base rather than maintaining separate systems. Multiple US states have subsequently enacted similar legislation.

8.3 EU Artificial Intelligence Act

The EU AI Act, agreed in 2024 and coming into full effect by 2026-2027, takes a risk-based approach to AI regulation, with different requirements for different risk levels. For marketing, the most relevant provisions concern:

Prohibited practices, including AI systems that deploy subliminal techniques beyond a person's consciousness to distort behaviour in a manner that causes harm, or that exploit vulnerabilities of specific groups. Real-time biometric identification for commercial purposes in public spaces is prohibited with limited exceptions.

High-risk AI systems (which may include some personalisation applications) subject to requirements for high-quality training data, transparency documentation, human oversight, and accuracy standards.

Transparency requirements for AI systems that interact with consumers - users must be informed when they are interacting with an AI system.

8.4 Federal Trade Commission (USA)

The FTC has increasingly focused on data collection practices, algorithmic harms, and deceptive advertising. Its commercial surveillance rulemaking process, initiated in 2022, examines whether comprehensive rules limiting the collection, use, and monetisation of consumer data are warranted.

The FTC has taken action against companies for deceptive claims about data practices, for targeting advertising to children in violation of COPPA (Children's Online Privacy Protection Act), and for using discriminatory ad targeting.

8.5 Digital Markets Act (EU)

The DMA, effective from 2023, designates the largest digital platforms (Google, Meta, Amazon, Apple, Microsoft, TikTok) as "gatekeepers" and imposes specific obligations relevant to advertising: requirements to enable advertisers to independently verify their ad delivery and performance; restrictions on combining data from different services without explicit user consent; interoperability requirements; and prohibition of self-preferencing.

IX. EMERGING TRENDS AND FUTURE DIRECTIONS

9.1 The Post-Cookie Era and Privacy-Preserving Advertising

The deprecation of third-party cookies - Google's delayed but advancing implementation for Chrome - is forcing a fundamental restructuring of the digital advertising ecosystem. Several technical approaches are being developed as alternatives:

Federated Learning of Cohorts (FLoC) and Privacy Sandbox: Google's proposed replacement for third-party cookies involves doing interest-based grouping on-device, sharing only cohort membership (not individual identity) with advertisers. This approach has faced criticism from privacy advocates (still enables fingerprinting) and publishers (reduces targeting precision).

First-Party Identity Solutions: Brands are investing heavily in building direct consumer relationships - email lists, loyalty programmes, account-based experiences - that provide first-party data not dependent on third-party infrastructure.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

Contextual AI Revival: Sophisticated AI-powered contextual targeting - understanding the meaning, sentiment, and relevance of page content - is being positioned as a privacy-respecting alternative to user-level targeting.
Clean Rooms: Privacy-preserving data collaboration environments where multiple parties can compute aggregate results across combined datasets without either party exposing individual-level data to the other.
Differential Privacy: A mathematical framework for adding calibrated noise to data to protect individual privacy while preserving aggregate statistical utility. Being applied to ad measurement and audience modelling.

9.2 Generative AI and Creative Personalisation

Generative AI is enabling a new frontier of creative personalisation - not just selecting among existing assets, but generating unique creative content for each impression:
Dynamic Copy Generation: LLMs generate unique ad headlines, body text, and calls-to-action calibrated to individual user profiles, going far beyond template-based personalisation.
Personalised Visual Creative: Diffusion models generate or modify images to incorporate individual-specific elements - a consumer's name, local imagery, culturally relevant visual references.
Interactive Ad Experiences: LLM-powered conversational interfaces within ad units enable consumers to ask questions about products, receive personalised recommendations, and progress through purchase journeys in natural conversation.
Automated Creative Testing at Scale: Generative AI enables the creation of thousands of creative variants simultaneously, with AI-driven testing systems identifying optimal combinations for different audience segments in real time.

9.3 Conversational Commerce and AI Assistants

LLM-powered conversational AI is creating new marketing touchpoints that blur the boundary between assistance and advertising:
AI shopping assistants - deployed on retail websites, messaging platforms, and increasingly as standalone products - provide personalised product guidance in natural conversation, combining informational assistance with commercial recommendation.
The integration of advertising into AI assistant responses (already being explored by Google and Microsoft in their AI search products) represents a fundamentally new context for commercial communication - one where the boundary between organic advice and paid promotion must be carefully managed.
Voice-activated commerce - AI assistants in smart speakers, cars, and wearables - is an emerging channel where personalised marketing operates in audio rather than visual formats, requiring entirely new creative and attribution approaches.

9.4 AI in Influencer Marketing and Creator Economy

AI is transforming influencer marketing - connecting brands with content creators - through:
AI-Powered Creator Discovery: ML systems analyse creator content, engagement patterns, audience demographics, and brand safety to identify optimal influencer partnerships.
Performance Prediction: Models predict the likely reach, engagement, and conversion impact of influencer campaigns before investment.
Virtual Influencers: AI-generated digital personas with their own aesthetics and personalities are increasingly used for brand endorsement, particularly in fashion and entertainment.
AI-Assisted Content Creation: Creators using AI tools to produce more content more efficiently, with implications for authenticity and disclosure.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

9.5 Emotional AI and Affective Computing

Research into emotional AI - systems capable of detecting and responding to consumer emotional states - is advancing rapidly, with profound commercial and ethical implications:

Sentiment Detection in Voice: Voice AI can infer emotional state from speech patterns, enabling real-time adaptation of conversational marketing to consumer mood.

Facial Expression Analysis: Computer vision systems can infer emotional response to ad creative from facial expressions captured via device cameras (with consent), enabling real-time creative optimisation.

Physiological Signal Analysis: Wearable devices collecting heart rate, skin conductance, and other biometric data enable inferring emotional engagement with content.

The use of such systems in advertising raises serious ethical concerns about the exploitation of emotional vulnerability and the appropriate scope of consumer consent.

9.6 The Metaverse and Immersive Commerce

Extended reality environments - augmented reality, virtual reality, and mixed reality - represent an emerging frontier for personalised marketing:

AR Try-On Experiences: AI-powered augmented reality enables consumers to virtually try on clothing, furniture, makeup, and other products, dramatically reducing purchase uncertainty.

In-World Advertising: Virtual environments enable the placement of contextually relevant virtual ads within immersive experiences, opening new inventory for personalised marketing.

Avatar-Based Personalisation: In virtual worlds, users' avatars and behavioural patterns provide rich signals for personalisation of virtual goods and experiences.

X. TOWARD RESPONSIBLE AI MARKETING: A FRAMEWORK

Given the documented benefits and risks, a framework for responsible AI marketing should encompass:

10.1 Privacy by Design

Privacy considerations should be embedded in AI marketing system design from the outset, not added as compliance afterthoughts. This means minimising data collection to what is genuinely necessary, implementing strong access controls and retention limits, using privacy-preserving techniques where possible, and proactively communicating with consumers about data practices.

10.2 Transparency and Consumer Control

Consumers should be able to understand why they are seeing particular ads, to access information about what data underpins targeting decisions, and to exercise meaningful control over their data and their ad experiences. This requires investment in consumer-facing transparency tools - ad preference managers, data portability, and clear consent mechanisms - that go beyond minimum regulatory compliance.

10.3 Algorithmic Auditing

AI personalisation systems should be subject to regular independent audit for discriminatory outcomes, manipulation of vulnerable populations, and compliance with stated policies. These audits should be required by regulation and results made available to regulators if not to the public.

10.4 Ethical Objective Setting

The objectives for which AI marketing systems are optimised have profound effects on outcomes. Systems optimised purely for short-term engagement or conversion maximisation will tend toward manipulation and exploitation. Incorporating longer-term consumer welfare, trust, and satisfaction into optimisation objectives - and measuring these outcomes - can align commercial incentives with ethical practice.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 5, April 2026 Impact Factor: 8.2

10.5 Industry Self-Regulation and Standards

Industry bodies - IAB, World Federation of Advertisers, Partnership on AI - have roles to play in establishing and enforcing standards for responsible personalised advertising that supplement regulatory requirements. Effective self-regulation requires credible enforcement mechanisms, not merely voluntary principles.

XI. CONCLUSION

Artificial intelligence has fundamentally transformed personalised advertising and marketing, enabling a level of individual relevance, operational efficiency, and commercial effectiveness that was impossible in any previous era. The technologies involved - from classical machine learning to deep neural networks, from NLP to generative AI - continue to advance rapidly, and their marketing applications continue to proliferate across channels, industries, and geographies.

The commercial case for AI-driven personalisation is compelling and well-documented. Brands that deploy these capabilities effectively achieve substantially superior acquisition, conversion, retention, and revenue outcomes. For consumers, well-implemented personalisation genuinely improves the relevance of commercial communication and facilitates product discovery.

REFERENCES

[1]. Vaswani, A., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
[2]. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
[3]. McKinsey & Company (2021). *The value of getting personalisation right - or wrong - is multiplying.* McKinsey Global Institute.
[4]. McKinsey & Company (2023). *The economic potential of generative AI: The next productivity frontier.*
[5]. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning.* MIT Press.
[6]. Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power.* PublicAffairs.
[7]. Matz, S. C., Kosinski, M., Nave, G., & Stillwell, D. J. (2017). Psychological targeting as an effective approach to digital mass persuasion. *PNAS*, 114(48), 12714–12719.
[8]. European Parliament and Council (2018). *General Data Protection Regulation (GDPR).* Regulation (EU) 2016/679.
[9]. European Parliament and Council (2024). *Artificial Intelligence Act.* Regulation (EU) 2024/1689.
[10]. European Parliament and Council (2022). *Digital Markets Act.* Regulation (EU) 2022/1925.
[11]. California Legislature (2018). *California Consumer Privacy Act (CCPA).* AB-375.
[12]. Federal Trade Commission (2022). *Commercial Surveillance and Data Security: Advance Notice of Proposed Rulemaking.*
[13]. Angwin, J., Parris Jr., T., & Mattu, S. (2016). Facebook Lets Advertisers Exclude Users by Race. *ProPublica.*
[14]. IAB & PwC (2023). *Internet Advertising Revenue Report: Full Year 2022.*
[15]. Ricci, F., Rokach, L., & Shapira, B. (Eds.) (2022). *Recommender Systems Handbook* (3rd ed.). Springer.
[16]. Salesforce Research (2022). *State of the Connected Customer* (5th ed.).
[17]. Netflix Technology Blog (2022). *Artwork Personalisation at Netflix.*
[18]. Covington, P., Adams, J., & Sargin, E. (2016). Deep neural networks for YouTube recommendations. *Proceedings of the 10th ACM Conference on Recommender Systems.*
[19]. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.
[20]. Dwork, C., et al. (2006). Calibrating noise to sensitivity in private data analysis. *Theory of Cryptography Conference.*