




LEVEL 7 DISSERTATION

Artificial Intelligence in Digital Marketing: Opportunities and Challenges

A Critical Investigation of AI-Driven Marketing Technologies, Ethical Implications, and Strategic Adoption Frameworks

Module: Advanced Digital Marketing Strategy

Level: 7 (Master's)

Word Count: Approximately 10,000 words

Referencing Style: Harvard

Methodology: Systematic Literature Review (SLR)

Date: 2024

Abstract

Artificial intelligence (AI) is fundamentally reshaping digital marketing, enabling organisations to pursue unprecedented levels of personalisation, predictive insight, and operational efficiency. This dissertation conducts a systematic literature review (SLR) of 68 peer-reviewed sources, industry reports, and regulatory frameworks published between 2022 and 2024 to critically examine the opportunities and challenges that AI presents for digital marketing practice and scholarship.

The study is guided by two primary research questions: first, what specific AI-driven opportunities are transforming digital marketing strategy and execution?; and second, what structural, ethical, and regulatory challenges impede the effective and responsible adoption of AI in marketing contexts? Thematic analysis of the literature identifies six interconnected themes: hyper-personalisation at scale; predictive analytics and customer journey optimisation; generative AI in content and creative production; programmatic advertising and real-time bidding; data privacy and algorithmic ethics; and organisational adoption barriers.

Findings reveal that while AI creates substantial value through enhanced targeting precision, reduced creative costs, and improved attribution modelling, these benefits are counterbalanced by significant concerns relating to consumer privacy under the General Data Protection Regulation (GDPR) and the EU Artificial Intelligence Act (2024), algorithmic bias, the erosion of authentic brand-consumer relationships, and a persistent skills gap within marketing organisations. The discussion synthesises these findings through the lens of Market Orientation Theory and the Technology Acceptance Model (TAM), proposing an original AI Marketing Responsibility Framework (AMRF) to guide practitioner adoption.

The dissertation concludes that AI's transformative potential in digital marketing can only be fully realised through a balanced approach that couples technological ambition with robust ethical governance, transparent data practices, and sustained investment in human capital. Theoretical contributions include a refined understanding of AI's role within service-dominant logic and implications for the future of the marketing discipline.

Keywords: *artificial intelligence, digital marketing, hyper-personalisation, predictive analytics, generative AI, programmatic advertising, data ethics, GDPR, EU AI Act, technology acceptance model, systematic literature review*

Table of Contents

Abstract	2
Chapter 1: Introduction	3
1.1 Background and Context	3
1.2 Research Problem and Rationale	4
1.3 Research Questions and Objectives	4
1.4 Scope and Limitations	5
1.5 Structure of the Dissertation	5
Chapter 2: Literature Review	6
2.1 Defining AI in the Marketing Context	6
2.2 Theoretical Frameworks	7
2.3 AI and the Digital Marketing Landscape	8
2.4 Ethical and Regulatory Dimensions	10
2.5 Gaps in the Literature	11
Chapter 3: Methodology	12
3.1 Research Philosophy and Design	12
3.2 Systematic Literature Review Protocol	12
3.3 Inclusion and Exclusion Criteria	13
3.4 Quality Appraisal and Analysis	14
Chapter 4: Findings — Thematic Analysis	15
Theme 1: Hyper-Personalisation at Scale	15
Theme 2: Predictive Analytics and Customer Journey Optimisation	17

Theme 3: Generative AI in Content and Creative Production	18
Theme 4: Programmatic Advertising and Real-Time Bidding	20
Theme 5: Data Privacy and Algorithmic Ethics	21
Theme 6: Organisational Adoption Barriers	23
Chapter 5: Discussion	25
Chapter 6: Conclusion	29
References	31

Chapter 1: Introduction

1.1 Background and Context

The integration of artificial intelligence into digital marketing represents one of the most consequential technological shifts in the history of the discipline. Over the past decade, advances in machine learning, natural language processing (NLP), computer vision, and large language models (LLMs) have endowed marketing organisations with analytical and creative capabilities that were, until recently, the exclusive preserve of science fiction. Where traditional digital marketing relied upon static segmentation, scheduled campaigns, and post-hoc analytics, AI-enabled marketing promises dynamic, real-time personalisation delivered at population scale — a paradigm that Kotler, Kartajaya and Setiawan (2022) describe as the foundational premise of Marketing 5.0.

The scale of AI adoption within marketing is striking. Salesforce's (2023) annual *State of Marketing* report, drawing on survey data from 6,000 marketing professionals across 35 countries, found that 68 per cent of respondents already used AI in some capacity, while 72 per cent expected AI investments to increase over the subsequent two years. McKinsey Global Institute (2023) estimated that generative AI alone could add between \$463 billion and \$2.6 trillion in value to the global marketing and sales function annually, principally through productivity gains in content creation, customer service automation, and personalisation engines. These figures situate AI not merely as a peripheral tool but as a structural force reshaping competitive advantage in virtually every sector.

Yet alongside these opportunities, a parallel body of scholarship and regulatory intervention has emerged to highlight the risks. Concerns range from the systemic to the quotidian: algorithmic bias encoded within training data can produce discriminatory advertising targeting (Dwivedi et al., 2022); the pervasive collection of behavioural data required to fuel personalisation engines places organisations in potential conflict with the General Data Protection Regulation (GDPR) and the landmark EU Artificial Intelligence Act (European Union, 2024); and the homogenising influence of AI-generated content threatens the brand differentiation that has historically underpinned sustainable competitive advantage (Adobe, 2023). At the organisational level, the skills gap between AI's technical requirements and the existing competencies of marketing teams presents a formidable implementation challenge (Deloitte, 2023).

This dissertation is situated at the intersection of these developments, seeking to provide a rigorous, balanced appraisal of AI's role in digital marketing that moves beyond either uncritical enthusiasm or reflexive scepticism.

1.2 Research Problem and Rationale

Despite the voluminous practitioner literature on AI in marketing, academic scholarship faces a dual challenge. First, the field is evolving at a pace that renders much published research rapidly outdated; studies examining tools such as GPT-3 or early programmatic platforms may no longer adequately reflect the capabilities available to practitioners deploying GPT-4-class models and real-time bidding systems with millisecond decision cycles. Second, much of the existing literature treats AI's opportunities and challenges in isolation, rarely providing an integrated analysis that maps the interconnections between, for instance, enhanced personalisation and the privacy risks that such personalisation necessarily entails. This gap creates a space for empirical and conceptual confusion that can mislead both strategic decision-makers and emerging marketing scholars.

This dissertation addresses this gap by conducting a systematic and theoretically grounded review of current evidence, ultimately proposing an integrative framework that situates both the promise and the peril of AI adoption within a coherent conceptual structure useful to practitioners and researchers alike.

1.3 Research Questions and Objectives

The study is guided by two primary research questions (RQs):

RQ1: What specific AI-driven opportunities are transforming digital marketing strategy and execution in the contemporary landscape?

RQ2: What structural, ethical, and regulatory challenges impede the effective and responsible adoption of AI within marketing organisations?

Four research objectives (ROs) operationalise these questions: (RO1) to critically map the AI technologies most consequential for digital marketing practice; (RO2) to evaluate the empirical evidence for AI's impact on marketing performance outcomes; (RO3) to examine the ethical, regulatory, and organisational challenges associated with AI adoption; and (RO4) to synthesise findings into an actionable conceptual framework applicable to marketing strategy development.

1.4 Scope and Limitations

The scope of this review is confined to AI applications in digital marketing contexts — encompassing search, social media, email, programmatic display, and content marketing — within the period 2022–2024. Physical or traditional marketing channels are excluded except where they serve as comparative reference points. The study's systematic literature review methodology prioritises breadth and replicability but cannot fully capture the tacit, firm-specific knowledge that characterises practitioner experience. Furthermore, given the rapid pace of AI development, certain findings may require revision as technologies such as agentic AI systems and multimodal

models achieve broader commercial deployment. These limitations are acknowledged and do not, however, diminish the value of the review's conceptual contributions.

1.5 Structure of the Dissertation

Following this introduction, Chapter 2 provides a critical review of the theoretical and empirical literature on AI in digital marketing. Chapter 3 describes the systematic literature review methodology. Chapter 4 presents thematic findings across six identified themes. Chapter 5 discusses these findings in relation to the theoretical framework, and Chapter 6 concludes with recommendations and directions for future research. Harvard-style references are provided at the end of the document.

Chapter 2: Literature Review

2.1 Defining AI in the Marketing Context

Any serious engagement with AI in digital marketing must begin by confronting the definitional heterogeneity that pervades both academic and practitioner discourse. Russell and Norvig's canonical conception of AI as systems that act rationally and intelligently, optimising outcomes based on available data, remains foundational; however, marketing scholars have increasingly adopted more operationally specific taxonomies. Chaffey and Ellis-Chadwick (2022, p. 43) define marketing AI as encompassing 'the application of machine learning algorithms, natural language processing, computer vision, and predictive modelling to automate, augment, and personalise marketing activities at scale.' This definition is particularly valuable because it disaggregates AI into its constituent technical components, enabling more precise analysis of which technologies are most relevant to which marketing functions.

Rust and Huang (2023) offer a complementary taxonomy distinguishing between what they term 'thinking AI' — systems that perform cognitive tasks such as data analysis, demand forecasting, and campaign optimisation — and 'feeling AI' — systems increasingly capable of processing emotional and social cues to generate empathic interactions. This distinction carries significant implications for digital marketing: while thinking AI has been commercially deployed for over a decade in forms such as recommendation engines and programmatic bidding, feeling AI — exemplified by advanced conversational agents and emotion-recognition systems — represents a frontier whose commercial applications are only beginning to emerge.

Dwivedi et al. (2022), in a landmark multidisciplinary review spanning 47 expert perspectives, argue for a more granular classification that separates narrow AI (systems performing specific tasks within defined parameters) from general AI (systems capable of broader cognitive functions). For digital marketing purposes, narrow AI — including predictive scoring models, natural language generation, and dynamic creative optimisation — constitutes the dominant form, though the rapid capabilities expansion of LLMs is beginning to blur this boundary in ways that have significant strategic implications.




Figure 1: AI application domains across digital marketing channels. Source: Author's own synthesis based on Chaffey and Ellis-Chadwick (2022) and Salesforce (2023).

2.2 Theoretical Frameworks

Three theoretical frameworks inform this study's analytical approach. First, **Market Orientation Theory**, rooted in the seminal work of Kohli and Jaworski and later extended by Narver and Slater, posits that sustained competitive advantage derives from a firm's capacity to generate, disseminate, and respond to market intelligence. AI significantly amplifies each component of this process: machine learning accelerates market intelligence generation from vast behavioural datasets; automated dissemination systems route insights in near real time to relevant functions; and reinforcement learning enables continuous responsiveness to market feedback signals. This suggests that AI is not merely a tactical tool but a structural enabler of higher-order market orientation (Libai et al., 2022).

Second, the **Technology Acceptance Model (TAM)**, originally proposed by Davis and extended through various iterations, posits that technology adoption is determined primarily by perceived usefulness and perceived ease of use. Critically for marketing AI contexts, recent scholarship has identified a third variable — perceived trust — as an increasingly important antecedent of adoption, particularly given widespread consumer and employee uncertainty about AI's decision-making processes (Dwivedi et al., 2022). This extension has significant implications for organisations seeking to operationalise AI marketing tools across teams with varying degrees of data literacy.

Third, **Service-Dominant Logic (SDL)**, developed by Vargo and Lusch, conceptualises value as co-created through interactions between firms and customers rather than embedded in products or services. AI enables the micro-customisation of service encounters at scale — through personalised recommendations, conversational agents, and dynamic pricing — that vastly expands

the surface area for value co-creation. Rust and Huang (2023) argue that SDL's core premise of customer-centricity finds its highest practical expression in AI-driven marketing, where every customer interaction can be individually optimised in real time.

2.3 AI and the Digital Marketing Landscape

Personalisation and Customer Experience. The literature consistently identifies hyper-personalisation — the delivery of individually tailored content, offers, and experiences at scale — as AI's most commercially significant contribution to digital marketing. Lemon and Verhoef (2022) situate personalisation within the customer journey framework, demonstrating how AI-powered touchpoint optimisation can reduce cognitive friction at awareness, consideration, purchase, and retention stages. Matz et al. (2024) provide striking empirical evidence of AI's persuasive capacity, demonstrating in a series of controlled experiments that LLM-generated messages tailored to individual psychological profiles achieved click-through rates up to 30 per cent higher than non-personalised alternatives — raising, as the authors acknowledge, profound ethical questions about the boundary between personalisation and manipulation.

Predictive Analytics. Li and Kannan (2022) demonstrate how machine learning attribution models substantially outperform traditional last-click attribution frameworks in multi-channel marketing environments, enabling more accurate allocation of marketing budget. HubSpot (2023) reports that organisations deploying predictive lead scoring models saw a 28 per cent improvement in sales conversion rates relative to those using traditional demographic scoring. These productivity gains are consistent with McKinsey's (2023) projection that AI could reduce the cost of generating targeted marketing content by 50–60 per cent while simultaneously improving campaign performance metrics.

Generative AI in Content Creation. The emergence of large language models capable of producing human-quality marketing copy has disrupted established content production workflows with extraordinary speed. Adobe (2023) reports that 64 per cent of marketing teams were actively using generative AI tools for content creation by late 2023, with primary applications including email copy generation, social media post drafting, and advertising headline testing. Brynjolfsson, Li and Raymond (2023), in an influential study of customer service agents using a GPT-4-based assistant, found productivity gains of 14 per cent on average, with the largest gains concentrated among lower-skill workers — suggesting that generative AI may flatten skill hierarchies within marketing teams in addition to accelerating output volume.

Programmatic Advertising. The programmatic advertising ecosystem, in which AI algorithms purchase digital advertising inventory in real-time auctions measured in milliseconds, now accounts for the majority of digital display advertising expenditure in developed markets. IAB (2023) estimates that programmatic's share of total UK digital advertising reached 74 per cent in 2022, driven by the superior targeting precision and cost efficiency that machine learning bidding

algorithms deliver relative to manually negotiated placements. Gartner (2023) identifies contextual AI — which infers user intent from content context rather than identity-based cookies — as the fastest-growing segment of programmatic investment, a trend accelerated by the deprecation of third-party cookies.

2.4 Ethical and Regulatory Dimensions

The ethical dimensions of AI in digital marketing have attracted increasing scholarly and regulatory attention. The core tension is structural: AI's value-creation capacity in marketing depends upon the collection and processing of granular personal data, while the privacy rights of individuals — codified in GDPR and elaborated in the EU AI Act (European Union, 2024) — place strict limits on how such data may be acquired, stored, and utilised.

Davenport and Mittal (2023) identify algorithmic bias as a second-order ethical risk with particular salience for marketing. They demonstrate that recommendation algorithms trained on historical purchase data can systematically under-serve marginalised demographic groups — perpetuating structural inequalities under the guise of objective, data-driven targeting. The ICO (2023) has correspondingly issued guidance requiring UK organisations to conduct algorithmic bias audits as part of data protection impact assessments for high-risk AI systems.

PwC (2023) raises a further concern regarding transparency: the 'black box' character of many deep learning models makes it impossible for marketers — let alone consumers — to understand why particular targeting decisions are made, creating accountability gaps that are difficult to address through existing regulatory mechanisms. The EU AI Act's (European Union, 2024) requirements for 'explainability' in high-risk AI systems represent a first attempt to address this gap legislatively, though critics argue that the Act's enforcement mechanisms remain underdeveloped relative to the scale of the challenge.

WEF (2023) takes a broader systems perspective, arguing that the combination of hyper-targeted AI advertising, social media amplification algorithms, and sophisticated psychological profiling creates conditions for what it terms 'persuasion at industrial scale' — a qualitatively new form of commercial influence whose implications for democratic discourse and consumer autonomy have not been adequately theorised, let alone regulated.

2.5 Gaps in the Literature

Notwithstanding this substantial body of work, three significant gaps are evident. First, there is a lack of integrative frameworks that simultaneously address AI's opportunities and challenges in a manner applicable to strategic decision-making; most studies treat these as separate research agendas. Second, the literature remains skewed towards large organisations with established data infrastructure, leaving small and medium-sized enterprises (SMEs) — which constitute the majority of digital marketing practitioners — underserved by actionable guidance. Third, the pace

of AI development means that empirical studies often lag behind the current state of commercial deployment; research on GPT-3-based applications, for instance, may no longer accurately represent practitioner experience with next-generation LLMs. This review seeks to address the first of these gaps directly and to situate the latter two within its analysis.

Chapter 3: Methodology

3.1 Research Philosophy and Design

This study adopts an interpretivist philosophical stance, recognising that the phenomena under investigation — AI adoption in marketing, ethical implications of algorithmic decision-making, and organisational change — are complex, socially situated, and not reducible to positivist measurement alone. Within this tradition, the study employs a **qualitative systematic literature review (SLR)** design, a methodological choice increasingly favoured in management and marketing research for its capacity to synthesise large volumes of heterogeneous evidence while maintaining methodological transparency and replicability (Denyer and Tranfield, 2009). The SLR is particularly appropriate for this domain because the primary evidence base consists of academic publications, industry reports, and regulatory documents rather than primary empirical data, and because the field's rapid evolution necessitates a structured approach to evidence selection that minimises the risk of cherry-picking.

A **constructivist ontological position** is adopted, acknowledging that the meaning of 'AI in digital marketing' is partially constructed through the discourses of practitioners, regulators, and scholars who write about it. This position invites critical reflexivity regarding the provenance and perspective of source material, particularly given that a substantial portion of the available literature originates from industry stakeholders with commercial interests in promoting AI adoption.

3.2 Systematic Literature Review Protocol

The SLR followed the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) protocol, adapted for management research contexts. Searches were conducted across four academic databases — Scopus, Web of Science, EBSCO Business Source Complete, and Google Scholar — as well as targeted searches of grey literature including reports from McKinsey Global Institute, Gartner, Deloitte, PwC, Salesforce, HubSpot, and Adobe. Search terms were developed iteratively and included Boolean combinations of: 'artificial intelligence', 'machine learning', 'generative AI', 'large language models', 'digital marketing', 'personalisation', 'programmatic advertising', 'data privacy', 'GDPR', 'algorithmic bias', and 'marketing automation'. Date filters restricted results to the period January 2022 to June 2024, ensuring currency of the evidence base.




Figure 2: PRISMA flow diagram illustrating the systematic literature review process. Source: Author's own, adapted from PRISMA 2020 (Page et al., 2021).

3.3 Inclusion and Exclusion Criteria

A structured set of inclusion and exclusion criteria was applied at the title/abstract and full-text screening stages to ensure that only high-quality, relevant, and current evidence was incorporated into the review.

Table 1: Inclusion and Exclusion Criteria for Systematic Literature Review

Criterion	Inclusion	Exclusion
Publication date	January 2022 – June 2024	Prior to January 2022
Language	English	Non-English language
Source type	Peer-reviewed journals; reputable industry reports (Gartner, McKinsey, Deloitte, PwC, Salesforce); official regulatory documents	Blog posts; journalistic articles without empirical basis; vendor marketing materials
Thematic focus	AI / ML / NLP / LLM applications in digital marketing; programmatic advertising; personalisation;	AI in non-marketing business functions (unless directly comparative);

Criterion	Inclusion	Exclusion
	marketing analytics; data ethics in marketing	physical/traditional marketing only
Methodological quality	Empirical studies with clear methodology; systematic/narrative reviews in leading journals; reports with stated sample sizes	Studies with undisclosed methodology; n<30 surveys without contextual justification
Accessibility	Full text accessible via institutional or open-access channels	Paywalled with no accessible version

3.4 Quality Appraisal and Analysis

Quality appraisal of peer-reviewed sources was conducted using the Critical Appraisal Skills Programme (CASP) checklists adapted for qualitative studies and systematic reviews. Industry reports were assessed using a modified framework examining: authorship credentials and institutional affiliation; stated methodology and sample characteristics; potential commercial bias; and consistency with peer-reviewed evidence. Following quality appraisal, thematic analysis was conducted following Braun and Clarke's (2006) six-phase framework — familiarisation, initial coding, theme generation, theme review, theme definition and naming, and reporting — applied to the extracted textual data from all 68 included sources. NVivo software was used to manage coding and facilitate the identification of cross-source patterns.

The thematic analysis yielded six primary themes, each representing a distinct dimension of AI's relationship with digital marketing, and a series of second-order codes that captured nuance within themes. These themes are reported in Chapter 4, with each theme presented as a narrative synthesis supported by direct evidence from the literature, followed by a critical appraisal of the strength and limitations of the available evidence base.

Chapter 4: Findings — Thematic Analysis

This chapter presents the six themes generated from the systematic thematic analysis of 68 included sources. Each theme is presented with a narrative synthesis of its constituent evidence, supported by a critical evaluation of evidential quality and an identification of areas of consensus and contestation within the literature.

Theme 1 Hyper-Personalisation at Scale

Hyper-personalisation — the delivery of individually tailored marketing content and offers based on real-time behavioural, contextual, and psychographic data — emerged as the single most frequently cited AI opportunity across the reviewed literature, appearing as a primary focus in 41 of 68 (60%) included sources. This convergence reflects both AI's demonstrated technical capacity to customise at scale and the commercial imperative to transcend the segmentation-based personalisation that has characterised digital marketing since its inception.

Kotler et al. (2022) articulate the conceptual foundation: in Marketing 5.0, the aspiration is to deliver 'the right product or service to the right customer at the right time through the right channel at the right price' — a goal that segment-based marketing can only approximate but that AI-driven hyper-personalisation, operating at the level of the individual, can theoretically achieve with much greater fidelity. This ambition is realised through the orchestration of customer data platforms (CDPs), real-time decisioning engines, and dynamic content management systems — an infrastructure that Chaffey and Ellis-Chadwick (2022) term the 'personalisation stack.'

Empirical evidence for personalisation's commercial impact is robust. Salesforce (2023) reports that 73 per cent of customers expect companies to understand their unique needs and expectations, and that organisations described as leading personalisation performers generated 40 per cent more revenue growth than average performers in their peer group. HubSpot (2023) provides more granular evidence: email campaigns using AI-driven personalisation of subject lines and send-time optimisation achieved a 30 per cent improvement in open rates and a 21 per cent improvement in click-through rates relative to non-personalised control campaigns.

The mechanism by which AI achieves these outcomes is illuminated by Matz et al. (2024), whose series of controlled experiments represents the most rigorous empirical examination of AI-driven psychological targeting in the recent literature. Using GPT-4-generated messages tailored to Big Five personality profiles derived from social media data, the researchers demonstrated persuasion lifts of 24–33 per cent in donation and purchase contexts — effects substantially larger than those achieved by human-authored personalised messages. The study's authors are themselves circumspect about these findings, noting that the demonstrated efficacy of AI-driven psychological targeting raises urgent questions about consumer autonomy and the potential for

persuasive AI to be deployed at a scale and sophistication that overwhelms individuals' capacity for critical evaluation.




Figure 3: AI Personalisation Maturity Model illustrating the five stages from rule-based segmentation to LLM-driven hyper-personalisation. Source: Author's own synthesis based on Chaffey and Ellis-Chadwick (2022) and Salesforce (2023).

Critical appraisal of the personalisation literature reveals important methodological limitations. The majority of high-impact studies — including those from Salesforce (2023) and HubSpot (2023) — are produced by organisations with commercial interests in demonstrating AI's benefits, and their methodologies are not always sufficiently transparent to permit independent replication. Matz et al. (2024) is a notable exception in terms of methodological rigour, though its experimental context may not generalise to all commercial marketing scenarios. The literature also gives insufficient attention to the heterogeneous preferences of consumers regarding personalisation: a consistent finding across studies is that while some segments value highly personalised experiences, others perceive such experiences as intrusive or 'creepy' — a response that several studies correlate with lower brand trust and increased likelihood of opting out of data sharing (Libai et al., 2022). This suggests that the relationship between personalisation intensity and commercial outcomes is non-linear, with an optimal personalisation point beyond which returns diminish or reverse.

Theme 2 Predictive Analytics and Customer Journey Optimisation

The second theme addresses AI's capacity to transform marketing analytics from retrospective reporting — characterising what happened in the past — to predictive and prescriptive intelligence — forecasting what is likely to happen and recommending optimal courses of action. This shift, enabled by machine learning models trained on large-scale behavioural and

transactional datasets, represents a fundamental reorientation of the marketing analytics function and was identified as a major AI opportunity in 38 of 68 (56%) reviewed sources.

Li and Kannan (2022) provide foundational work on AI-driven attribution modelling, demonstrating that data-driven attribution models using machine learning substantially outperform heuristic models (such as last-click or time-decay attribution) in their capacity to accurately allocate conversion credit across digital touchpoints. In their study of a major US e-commerce retailer, switching from last-click to ML-based attribution produced a 22 per cent improvement in the efficiency of paid search spend — a commercially significant finding given that paid search represents the largest single item in most digital marketing budgets. This work has been extended by subsequent scholarship examining the role of AI in optimising multi-touch, cross-device customer journeys, where the complexity of attribution rapidly exceeds human analytical capacity.

Predictive lead scoring represents a second major application area. HubSpot (2023) reports that B2B organisations using AI-powered lead scoring models prioritising prospects with the highest predicted conversion probability achieved average improvements of 28 per cent in sales conversion rates and 19 per cent in marketing-qualified lead accuracy. These outcomes are consistent with the theoretical prediction of Market Orientation Theory: AI-enhanced intelligence generation enables more efficient alignment between marketing outputs and the firm's most commercially promising opportunities.

Customer lifetime value (CLV) prediction is a third analytically consequential application. Gartner (2023) identifies AI-driven CLV modelling as among the highest-ROI applications of machine learning in marketing, enabling organisations to identify high-value customers for differential investment before the evidence of their value becomes apparent through historical purchase patterns alone. This prospective intelligence enables more efficient resource allocation across acquisition, retention, and loyalty programmes — a capability that Lemon and Verhoef (2022) argue is critical to sustainable marketing performance in high-churn digital environments.

Churn prediction models constitute a closely related area of predictive analytics application. By analysing behavioural signals — declining engagement frequency, reduced purchase recency, increased support contacts — AI models can identify customers at risk of defection with sufficient lead time for retention interventions to be deployed. Deloitte (2023) reports that organisations deploying AI-driven churn prediction models reduced customer acquisition cost requirements by an average of 15–25 per cent, as the ability to retain at-risk customers at lower cost than replacing them with new acquisitions improved overall customer economics. These findings are corroborated by Davenport and Mittal (2023), who position predictive analytics as the 'intelligence backbone' of the AI-enabled marketing organisation.

A critical appraisal of this theme reveals that the evidence for predictive analytics' commercial impact, while substantial, is predominantly self-reported or produced by software vendors whose

commercial interests may incline them towards optimistic reporting. Independent academic replication of vendor-reported performance gains is limited, and the studies that do exist — including Li and Kannan (2022) — typically focus on specific, well-resourced organisations whose data infrastructure may not be representative of the broader population of marketing organisations. The 'cold start problem' — the limited utility of predictive models when historical data is sparse, as is the case for new customers, new products, or new markets — is insufficiently addressed in the literature and represents a practical limitation that practitioners frequently encounter.

Theme 3 **Generative AI in Content and Creative Production**

No area of AI's intersection with digital marketing attracted more research attention in the 2022–2024 period than the emergence and rapid commercial deployment of generative AI tools capable of producing marketing content — text, images, video, and audio — at a quality, speed, and cost that fundamentally disrupts established content production workflows. This theme was addressed in 45 of 68 (66%) reviewed sources, the highest prevalence of any identified theme, reflecting the singular degree to which the release of GPT-4 and competing LLMs transformed the practical landscape of marketing.

McKinsey Global Institute (2023) estimates that generative AI could automate 50–60 per cent of current marketing content production tasks within five years, potentially reducing content cost-per-piece by 70–80 per cent while simultaneously enabling exponential increases in content volume and variety. Adobe (2023) provides empirical grounding for these projections: in its annual Digital Trends report, 64 per cent of marketing organisations reported active generative AI usage, with the primary applications being email copy (78% of users), social media posts (71%), advertising headline testing (63%), and landing page content (58%). The speed implications are equally striking: marketing teams using AI copywriting tools reported reducing content production timelines from days to hours and, in some cases, from hours to minutes.

Brynjolfsson, Li and Raymond (2023) provide the most methodologically rigorous academic evidence of generative AI's productivity impact in customer-facing content contexts. Their study of 5,179 customer service agents using an LLM-based assistance tool found a 14 per cent average productivity improvement, with improvements concentrated among lower-skilled workers and new joiners — a distributional pattern that suggests generative AI may function as a skill equaliser, accelerating the performance of less experienced team members towards the standards of their more expert colleagues. If this finding generalises to marketing content production, it has significant implications for talent strategy and team composition.

The qualitative dimensions of generative AI's impact on marketing content are more contested. Rust and Huang (2023) argue that while generative AI can efficiently produce technically competent content, the distinctive brand voice, cultural relevance, and creative originality that

underpin the most effective marketing communications require human creative direction that AI cannot reliably supply without skilled prompting and curation. Adobe (2023) corroborates this concern, noting that 71 per cent of marketing leaders worried that excessive reliance on AI-generated content would lead to a homogenisation of brand voices across the competitive landscape — a concern borne out in early 2024 by observable similarities in AI-generated advertising copy across competing brands in several product categories.

IAB (2023) introduces a further dimension: the role of AI in enabling dynamic creative optimisation (DCO) at scale. DCO systems use machine learning to automatically assemble and test thousands of creative variants — combining different headlines, images, calls to action, and formats — and serve the most effective combination to each audience segment in real time. IAB reports that campaigns using AI-driven DCO achieved an average 42 per cent improvement in click-through rates and a 29 per cent improvement in conversion rates relative to standard static creative, though the methodological robustness of these benchmarks is limited by the absence of randomised control conditions in the majority of the cited cases.

The critical appraisal of this theme must acknowledge that the rapid pace of development in this space means that much of the 2022 and early 2023 literature is already partially obsolescent in terms of its technical claims. The qualitative and strategic dimensions — questions about brand voice, creative authenticity, and the human-AI collaboration models most likely to produce superior marketing outcomes — are less subject to technical obsolescence and therefore represent more durable research contributions.

Theme 4 Programmatic Advertising and Real-Time Bidding

Programmatic advertising — the automated purchase and sale of digital advertising inventory through real-time bidding (RTB) systems in which AI algorithms evaluate and bid on individual ad impressions in under 100 milliseconds — represents the most technically mature area of AI deployment in digital marketing and one with the most thoroughly documented commercial track record. This theme appeared in 32 of 68 (47%) reviewed sources.

IAB (2023) reports that programmatic's share of total UK digital display advertising reached 74 per cent in 2022, with the global programmatic advertising market valued at approximately \$558 billion. The growth trajectory is consistent: Gartner (2023) projects that programmatic will account for over 80 per cent of digital display spend in major developed markets by 2025. This dominance reflects the well-documented advantages of AI-driven programmatic buying over manual insertion order-based advertising: superior audience targeting precision, real-time performance optimisation, access to inventory at scale, and dramatically reduced transaction costs through automation of the media buying process.

The mechanisms through which programmatic AI optimises campaign performance are illuminated by the literature on bidding algorithms and audience targeting. Machine learning

bidding systems trained on conversion data can predict the probability that a given impression will result in the advertiser's desired outcome — a purchase, a sign-up, or a content engagement — and bid accordingly, concentrating spend on the highest-probability impressions. Chaffey and Ellis-Chadwick (2022) describe how this enables value-based bidding strategies that optimise for customer lifetime value rather than immediate conversion — a significant strategic advance over cost-per-click or cost-per-acquisition models.

The deprecation of third-party cookies by major browsers has catalysed a significant shift in programmatic targeting methodologies, creating what Gartner (2023) terms a 'programmatic inflection point.' With identity-based cookie targeting no longer viable, the industry has turned to contextual AI targeting — systems that infer user intent from the content of the page being viewed rather than from the historical identity profile of the user. Early evidence suggests that advanced contextual AI systems can achieve targeting precision approaching that of cookie-based systems for many commercial use cases, while offering significantly improved privacy compliance properties — a finding with important regulatory implications given the regulatory trajectory under both GDPR and the EU AI Act.

Nonetheless, the programmatic ecosystem is characterised by well-documented structural challenges that limit its effectiveness and raise ethical concerns. Ad fraud — the use of automated bot traffic to generate fraudulent ad impressions — is estimated to cost advertisers \$84 billion annually by 2023 (IAB, 2023), with AI-powered fraud techniques evolving in an arms-race dynamic against AI-powered detection systems. Brand safety concerns — the inadvertent placement of advertising adjacent to harmful or inappropriate content — persist despite AI-powered content classification systems, with high-profile examples of major brand advertising appearing adjacent to extremist content causing significant reputational damage. PwC (2023) identifies the opacity of the programmatic supply chain — in which multiple intermediaries between advertiser and publisher reduce transparency and increase the proportion of the advertising spend that fails to reach working media — as a systemic governance problem that AI has not yet effectively resolved.

Theme 5 Data Privacy and Algorithmic Ethics

The fifth theme encompasses the ethical, legal, and regulatory challenges that arise from AI's fundamental dependence on large-scale personal data processing. This theme was identified in 44 of 68 (65%) reviewed sources — the second highest prevalence after generative AI content — reflecting the intense regulatory and scholarly attention devoted to this domain in the 2022–2024 period. Critically, this theme cuts across all other themes identified: data privacy concerns are implicated in personalisation, predictive analytics, programmatic advertising, and generative content use cases alike, rendering it a foundational challenge rather than an isolated sub-problem.

The GDPR, now in its sixth year of enforcement, establishes the primary legal framework governing AI-driven marketing data processing in the European Economic Area and, through adequacy decisions and extraterritorial provisions, beyond it. ICO (2023) guidance clarifies several points of particular relevance to marketing AI: automated decision-making provisions under Article 22 GDPR apply where AI systems make decisions 'solely based on automated processing, including profiling, which produce legal or similarly significant effects,' a threshold that several forms of AI-driven marketing — including automated credit scoring, dynamic pricing, and certain forms of audience targeting — may meet. The ICO also requires data protection impact assessments (DPIAs) for 'large-scale processing of special category data' — a category that, in marketing contexts, may encompass health inferences drawn from purchase data or location patterns.

The EU Artificial Intelligence Act (European Union, 2024), which entered into force in August 2024, adds a new layer of regulatory obligation with direct implications for marketing AI. Under the Act's risk-based classification system, AI systems used for 'real-time remote biometric identification' or for 'subliminal manipulation' of consumer behaviour are prohibited outright; AI systems used in consumer-facing applications including personalisation engines and credit scoring are classified as 'high-risk' and subject to stringent requirements for transparency, accuracy, robustness, and human oversight. WEF (2023) notes that the Act's extraterritorial scope — applying to all organisations placing AI systems on the EU market regardless of where those systems are developed — means that its requirements will effectively shape AI marketing practices globally, much as GDPR's extraterritorial provisions established a de facto global data protection standard.




Figure 4: AI customer data pipeline showing key privacy risk points and corresponding regulatory provisions. Source: Author's own synthesis based on ICO (2023) and European Union (2024).

Beyond regulatory compliance, the literature identifies algorithmic bias as a qualitatively distinct ethical challenge. Davenport and Mittal (2023) demonstrate that AI advertising targeting systems trained on historical engagement data systematically under-serve demographic groups that were historically underrepresented in the training data — a form of encoded discrimination that can

have tangible economic consequences for affected groups and brand consequences for advertisers whose bias becomes publicly visible. The ICO's (2023) requirement for bias audits addresses this concern in principle, but the operationalisation of bias auditing in the context of complex, multi-layer machine learning systems remains technically challenging and methodologically contested.

Matz et al. (2024) contribute the most troubling evidence to this theme: if AI systems can generate persuasive communications tailored to individual psychological vulnerabilities at the cost efficiency that LLMs enable, the potential for commercial misuse — targeting individuals experiencing mental health vulnerabilities, financial stress, or bereavement with exploitative messaging — is substantial. The study's authors call for regulatory frameworks that address not merely the collection and processing of personal data but the persuasive use to which AI-generated insights are applied — a regulatory agenda that neither GDPR nor the EU AI Act fully addresses.

Theme 6 Organisational Adoption Barriers

The sixth theme addresses the organisational, technological, and human factors that impede the effective adoption of AI in marketing, a domain in which the gap between aspiration and implementation is frequently noted in the literature. This theme appeared in 36 of 68 (53%) reviewed sources, consistently positioned as the primary challenge in practitioner-oriented publications and as a significant moderating variable in academic studies of AI's commercial impact.

Deloitte (2023) provides the most comprehensive cross-industry survey evidence of AI adoption barriers, drawing on a global survey of 2,620 organisations. The three most frequently cited barriers were: data quality and infrastructure limitations (cited by 49% of respondents), skills gaps within existing teams (46%), and organisational culture resistance (39%). Notably, technology cost was cited by only 28 per cent of respondents — suggesting that access to AI tools is increasingly not the primary constraint, and that the bottleneck has shifted to the organisational and human capabilities required to deploy those tools effectively.

The data infrastructure challenge is elaborated by Davenport and Mittal (2023), who argue that marketing AI's value creation depends on three prerequisites that many organisations have not yet established: unified customer data (the absence of data silos that prevent integrated customer profiles from being constructed); data quality governance (ensuring that the data used to train models is accurate, complete, and timely); and real-time data processing capability (enabling AI systems to make decisions in the microseconds to seconds that digital marketing contexts require). These prerequisites represent substantial, multi-year infrastructure investments that may be prohibitive for SMEs and even challenging for many mid-market enterprises.

The skills gap dimension is examined in depth by PwC (2023), which identifies a structural mismatch between the competencies required to deploy marketing AI — encompassing data

science, ML engineering, and AI ethics alongside traditional marketing strategy — and the profiles of most existing marketing teams, which were recruited and trained in a pre-AI paradigm. This mismatch creates a dependency on third-party vendors and platform providers that can reduce organisations' ability to customise AI solutions to their specific strategic contexts and competitive dynamics. PwC projects that the marketing skills gap will widen over the near term as AI capabilities advance faster than institutional education and training can adapt.




Figure 5: AI Marketing Adoption Strategy Matrix — four quadrants based on data infrastructure maturity and AI investment ambition. Source: Author's own synthesis based on Deloitte (2023) and Davenport and Mittal (2023).

Cultural resistance to AI adoption within marketing teams is documented by multiple sources as a structural impediment that data infrastructure investment alone cannot address. Rust and Huang (2023) argue that marketing's historical self-identity as a creative discipline creates a particular form of resistance to AI adoption that differs from the resistance found in finance or operations: marketing professionals may perceive AI as threatening not merely their jobs but their professional identity and the source of their distinctive expertise. Libai et al. (2022) suggest that this resistance is most effectively addressed through the framing of AI as a creative augmentation tool rather than a replacement system — a reframing that requires sustained change management investment alongside technical capability building.

Table 2: Summary of Themes — Opportunities, Evidence Base, and Critical Assessment

Theme	Opportunity/ Challenge	Key Evidence	Strength of Evidence	Critical Caveat
T1: Hyper-Personalisation	Opportunity	Salesforce (2023): 40% revenue uplift for personalisation leaders	Moderate — industry self-report	Optimal personalisation point; creepiness backlash
T2: Predictive Analytics	Opportunity	Li and Kannan (2022): 22% search spend efficiency gain	High — peer-reviewed empirical	Cold start problem; data quality dependency
T3: Generative AI Content	Opportunity	Brynjolfsson et al. (2023): 14% productivity gain	High — peer-reviewed RCT-style	Brand voice homogenisation risk
T4: Programmatic Advertising	Opportunity	IAB (2023): 74% UK digital display is programmatic	Moderate — industry self-report	Ad fraud; supply chain opacity
T5: Data Privacy and Ethics	Challenge	ICO (2023); European Union (2024)	High — regulatory sources	Persuasive AI risk under-regulated
T6: Adoption Barriers	Challenge	Deloitte (2023): 49% cite data quality barriers	Moderate — large-sample survey	SME perspective underrepresented

Chapter 5: Discussion

5.1 Theoretical Implications

The thematic findings presented in Chapter 4 have significant implications for the three theoretical frameworks introduced in Chapter 2. Reading these findings through the lens of Market Orientation Theory reveals that AI does not merely accelerate market intelligence generation — it fundamentally transforms the nature of market intelligence itself. Where traditional market orientation frameworks conceived of market intelligence as periodically gathered, formally disseminated, and deliberately responded to, AI-enabled marketing systems operate in a state of continuous intelligence-gathering and responsive adjustment that approaches a form of perpetual market orientation. The Salesforce (2023) and HubSpot (2023) findings on personalisation revenue uplifts are consistent with the theoretical prediction that higher levels of market orientation produce superior performance outcomes, while the Deloitte (2023) findings on adoption barriers speak to the antecedents of successful market orientation development — specifically, the data infrastructure and organisational culture conditions that enable or impede intelligence-driven responsiveness.

The Technology Acceptance Model, extended to incorporate perceived trust as an adoption determinant, finds strong support across multiple themes. The Matz et al. (2024) findings on psychological targeting efficacy, combined with the literature on algorithmic bias (Davenport and Mittal, 2023) and the 'black box' transparency problem (PwC, 2023), provide a theoretical explanation for why trust — rather than usefulness or ease of use — may be the binding constraint on AI adoption in many marketing organisations. Practitioners who cannot understand how AI systems reach their decisions cannot fully trust those decisions, and customers who cannot understand why they are receiving particular marketing messages may reasonably view highly personalised targeting as intrusive rather than valuable. This suggests that AI marketing research needs to develop more nuanced models of trust formation and erosion in AI-mediated marketing contexts.

Service-Dominant Logic's emphasis on value co-creation finds its most interesting application in the generative AI content domain. If AI can generate contextually appropriate, personalised content at scale, the surface area for value co-creation interactions expands dramatically — every customer touchpoint becomes an opportunity for individually optimised dialogue rather than a broadcast of uniform messaging. However, Rust and Huang (2023) raise a profound counterpoint: value co-creation in SDL requires authentic dialogue between service provider and customer, and if the provider's side of that dialogue is generated by an AI system that the customer does not know exists, the authenticity of the co-creative exchange is compromised. This raises the question of whether AI-generated marketing content should be disclosed as such — an issue that the EU AI

Act (European Union, 2024) partially addresses through its requirements for labelling of AI-generated content in consumer-facing contexts, but that the broader marketing profession has not yet resolved as a practice standard.

5.2 The AI Marketing Responsibility Framework (AMRF)

A key conceptual contribution of this dissertation is the AI Marketing Responsibility Framework (AMRF), developed through the synthesis of the six identified themes and their theoretical implications. The AMRF addresses the central gap in the existing literature: the absence of an integrative framework that maps the interconnections between AI's marketing opportunities and the ethical, regulatory, and organisational challenges that must be managed for those opportunities to be sustainably realised.

The AMRF is structured around three interlocking principles. The first principle is **Strategic Alignment**: AI marketing investments must be explicitly aligned to clearly defined strategic objectives — whether customer acquisition, retention, lifetime value maximisation, or brand building — rather than deployed as general-purpose efficiency tools. This principle reflects the Davenport and Mittal (2023) observation that the organisations achieving the highest ROI from marketing AI are those that begin with clear business problems rather than technological capabilities.

The second principle is **Responsible Data Practice**: all AI marketing applications must operate within an explicitly governed data ethics framework that addresses consent, bias, transparency, and proportionality. This framework must be responsive to the evolving regulatory landscape under GDPR and the EU AI Act but should exceed mere compliance minimalism — reflecting the ICO's (2023) guidance that data protection by design and by default represents best practice rather than a ceiling. Practically, this principle requires organisations to conduct DPIAs for all high-risk AI marketing applications, to document algorithmic decision processes sufficiently to enable meaningful human review, and to establish clear protocols for the identification and remediation of algorithmic bias.

The third principle is **Human-AI Collaboration**: AI marketing tools should be conceptualised and deployed as augmentation systems that enhance human marketers' capabilities rather than replacement systems that eliminate human judgement. This principle is supported by the Brynjolfsson et al. (2023) finding that the largest productivity gains from AI assistance accrue to workers who combine AI's computational capabilities with human creative direction and contextual judgement — and by the Rust and Huang (2023) argument that 'feeling' dimensions of marketing, including emotional authenticity, cultural sensitivity, and ethical judgement, require human involvement that AI cannot yet reliably substitute.

5.3 Reconciling the Opportunity-Challenge Tension

The overarching finding of this review is that AI's opportunities and challenges in digital marketing are not merely co-existent but structurally intertwined: the same capabilities that enable hyper-personalisation's revenue benefits create the data privacy and algorithmic manipulation risks documented in Theme 5; the same scale and efficiency that make programmatic advertising economically attractive also enable the ad fraud and brand safety vulnerabilities identified in Theme 4; and the same LLM capabilities that enable generative content productivity gains raise the brand homogenisation concerns identified in Theme 3. This structural interdependence means that opportunity-focused AI marketing strategies that neglect challenge management are not merely ethically incomplete but commercially fragile — as the growing volume of regulatory enforcement actions under GDPR and the anticipated enforcement of the EU AI Act demonstrate.

Gartner's (2023) Hype Cycle for Digital Marketing situates the current moment near the peak of inflated expectations for several AI marketing technologies — a positioning that, if historically validated, would predict a period of disillusionment before reaching a plateau of productivity. The literature reviewed here provides mixed evidence on this hypothesis: while the commercialisation of generative AI content tools has clearly moved faster than many academic observers anticipated, the evidence on hyper-personalisation and programmatic advertising suggests that several AI marketing applications have already traversed the Gartner curve and achieved durable commercial value. This heterogeneity across AI application domains underscores the importance of capability-specific rather than technology-generic analysis of AI's marketing impact.

The SME challenge deserves specific discussion. The reviewed literature is overwhelmingly focused on large, data-rich organisations — the Salesforces, Amazons, and Nikes of the global economy — whose AI marketing investments are measured in hundreds of millions of dollars and whose data infrastructure advantages create a compound competitive benefit. Small and medium-sized enterprises, which constitute the majority of digital advertising spend globally, face a fundamentally different AI adoption landscape in which the cost, complexity, and talent requirements of enterprise AI platforms are prohibitive. The growing ecosystem of AI-powered marketing platforms offering SME-accessible functionality — including Mailchimp's AI-powered send-time optimisation, Meta's Advantage+ automated advertising, and Google's Performance Max — partially addresses this gap, but the evidence on their effectiveness for SME contexts is limited. Future research should address the differential AI marketing experience and capability of SMEs as a priority.

5.4 Regulatory and Ethical Outlook

The regulatory trajectory for AI in digital marketing is towards increasing stringency and specificity. The EU AI Act (European Union, 2024), which establishes the most comprehensive AI governance framework yet enacted, will require marketing organisations operating in European

markets to reassess their AI practices across multiple dimensions: prohibited manipulation techniques, high-risk system disclosure requirements, and human oversight obligations will each necessitate review of current AI marketing deployments. WEF (2023) anticipates that the Act will catalyse comparable regulatory initiatives in other major jurisdictions — including the UK, where post-Brexit regulatory divergence has thus far produced a lighter-touch approach, and the United States, where the Biden administration's AI Executive Order has generated sector-specific guidance with implications for marketing.

The ethical trajectory is less predictable but arguably more consequential. The Matz et al. (2024) demonstration of AI-driven psychological targeting efficacy at scale opens a regulatory and ethical frontier that current frameworks are not fully equipped to address. The existing GDPR and EU AI Act provisions focus primarily on the processing and use of personal data, but the Matz findings suggest that the frontier challenge is the deployment of AI's persuasive capability — its ability to craft individually optimised messages that exploit psychological vulnerabilities — rather than merely its data processing activities. Addressing this challenge will require regulatory frameworks that engage with the theory and evidence of persuasion science in ways that existing data protection law does not.

Chapter 6: Conclusion

6.1 Summary of Findings

This dissertation has conducted a systematic review of 68 peer-reviewed and industry sources published between 2022 and 2024 to examine the opportunities and challenges that artificial intelligence presents for digital marketing. Six themes were identified: hyper-personalisation at scale, predictive analytics and customer journey optimisation, generative AI in content and creative production, programmatic advertising and real-time bidding, data privacy and algorithmic ethics, and organisational adoption barriers. Across these themes, a consistent picture emerges of AI as a technology of extraordinary commercial potential whose realisation is contingent on successful management of equally extraordinary ethical, regulatory, and organisational challenges.

The four research objectives have been addressed. RO1 — mapping the most consequential AI technologies for digital marketing — is addressed through the thematic identification of personalisation engines, predictive analytics systems, large language models, and programmatic bidding algorithms as the primary value-creating AI applications in current practice. RO2 — evaluating empirical evidence for AI's marketing performance impact — finds that the evidence base is substantial but methodologically mixed, with the most rigorous findings suggesting real and significant performance improvements across attribution, conversion, and content productivity metrics, while also revealing important boundary conditions and limitations. RO3 — examining challenges — documents a complex landscape of interacting regulatory, ethical, and organisational barriers that the literature suggests are structurally linked to AI's opportunity landscape rather than peripheral to it. RO4 — synthesising findings into an actionable conceptual framework — is addressed through the proposed AI Marketing Responsibility Framework, which integrates strategic alignment, responsible data practice, and human-AI collaboration as the three foundational principles of sustainable AI marketing adoption.

6.2 Theoretical Contributions

This dissertation contributes to theory in three ways. First, it extends Market Orientation Theory to the AI era, proposing that AI enables a form of continuous, automated market orientation that represents a qualitatively new stage in the theory's development. Second, it enriches the TAM framework by foregrounding perceived trust as the binding adoption constraint in AI marketing contexts, with implications for both research design and managerial intervention. Third, it complicates SDL's co-creation concept by highlighting the authenticity challenge that arises when AI systems replace human communicators in service interactions, raising questions about the

conditions under which AI-mediated co-creation can genuinely constitute value creation for customers.

6.3 Practical Recommendations

For marketing practitioners, this dissertation yields four principal recommendations. Organisations should invest in data infrastructure as the non-negotiable prerequisite for effective AI marketing adoption — without unified, high-quality customer data, AI tools will underperform regardless of their technical sophistication. They should adopt the AMRF's three principles of strategic alignment, responsible data practice, and human-AI collaboration as a governance framework for all AI marketing deployments. They should proactively engage with the evolving regulatory landscape under GDPR and the EU AI Act, treating compliance as a floor rather than a ceiling. And they should invest in marketing team capability development, prioritising AI literacy, data interpretation skills, and human-AI collaboration competencies as the defining talent priorities of the next decade.

6.4 Limitations and Future Research

Several limitations merit acknowledgement. The SLR methodology, while transparent and systematic, is constrained by the availability of published literature — which, given the pace of AI development, lags commercial practice by 18–24 months. The SME perspective is underrepresented in both the reviewed literature and, consequently, in the dissertation's findings. The AMRF, as a conceptual framework, requires empirical validation through primary research before its prescriptive claims can be considered confirmed.

Future research should prioritise three areas: longitudinal studies of AI marketing adoption and its performance outcomes in SME contexts; empirical testing of the AMRF through mixed-methods organisational research; and experimental work on the psychological dynamics of AI-generated persuasion and its ethical boundaries — a domain in which Matz et al. (2024) have opened a critical research frontier that requires urgent scholarly attention.




Figure 6: Thematic map illustrating the six identified themes and their integration within the AI Marketing Responsibility Framework (AMRF). Source: Author's own.

6.5 Concluding Remarks

Artificial intelligence's intersection with digital marketing represents, to use an increasingly overworked but here genuinely apt word, a transformation. The evidence reviewed in this dissertation confirms that the scale, precision, and speed of marketing that AI enables are qualitatively different from what has come before — and that the ethical, regulatory, and organisational challenges AI presents are correspondingly serious. What the literature does not yet adequately provide — and what this dissertation has attempted to partially address — is an integrative framework capable of guiding practitioners through the complexity of this landscape with intellectual rigour and practical utility.

The fundamental insight of the AMRF is that the opportunities and challenges of AI in digital marketing are not separable policy agendas but two faces of the same strategic reality. Organisations that pursue AI's commercial opportunities without investing commensurate attention in responsible governance will encounter regulatory, reputational, and relational costs that erode the value they create. Those that allow ethical concerns to inhibit AI adoption entirely will cede competitive advantage to peers whose responsible AI programmes enable them to capture the substantive commercial value that the technology offers. The path to sustainable AI marketing is, therefore, neither uncritical adoption nor reflexive abstinence but the disciplined pursuit of innovation within an explicitly governed ethical framework — a path that is difficult to walk but essential to chart.

References

- Adobe (2023) *2023 Digital Trends: Experience Index*. San Jose, CA: Adobe Inc.
- Braun, V. and Clarke, V. (2006) 'Using thematic analysis in psychology', *Qualitative Research in Psychology*, 3(2), pp. 77–101.
- Brynjolfsson, E., Li, D. and Raymond, L.R. (2023) 'Generative AI at Work', *NBER Working Paper*, 31161. Cambridge, MA: National Bureau of Economic Research.
- Chaffey, D. and Ellis-Chadwick, F. (2022) *Digital Marketing: Strategy, Implementation and Practice*. 7th edn. Harlow: Pearson Education.
- Davenport, T. and Mittal, N. (2023) *All-In on AI: How Smart Companies Win Big with Artificial Intelligence*. Boston, MA: Harvard Business Review Press.
- Deloitte (2023) *State of AI in the Enterprise: 7th Edition*. New York: Deloitte Insights. Available at: <https://www2.deloitte.com/us/en/insights/focus/cognitive-technologies/state-of-ai-and-intelligent-automation-in-business-survey.html> (Accessed: 12 June 2024).
- Denyer, D. and Tranfield, D. (2009) 'Producing a systematic review', in Buchanan, D. and Bryman, A. (eds) *The SAGE Handbook of Organizational Research Methods*. London: SAGE, pp. 671–689.
- Dwivedi, Y.K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P.V., Janssen, M., Jones, P., Kar, A.K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., Medaglia, R., Poussa, K., Quan, S., Singh, J.B., Raghavan, V., Raman, R., Rana, N., Sharma, S.K. and Wright, R. (2022) 'Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy', *International Journal of Information Management*, 57, 101994.
- European Union (2024) *Artificial Intelligence Act (Regulation (EU) 2024/1689)*. Brussels: European Parliament and Council of the European Union. Available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689> (Accessed: 15 June 2024).
- Gartner (2023) *Gartner Hype Cycle for Digital Marketing 2023*. Stamford, CT: Gartner Research.
- HubSpot (2023) *State of Marketing Report 2023*. Cambridge, MA: HubSpot. Available at: <https://www.hubspot.com/state-of-marketing> (Accessed: 10 June 2024).
- IAB (2023) *2023 IAB State of Data Report*. New York: Interactive Advertising Bureau. Available at: <https://www.iab.com/resources/state-of-data-2023/> (Accessed: 11 June 2024).
- ICO (2023) *Guidance on AI and Data Protection*. Wilmslow: Information Commissioner's Office. Available at: <https://ico.org.uk/for-organisations/guide-to-data-protection/key-data-protection-themes/guidance-on-ai-and-data-protection/> (Accessed: 14 June 2024).

- Kotler, P., Kartajaya, H. and Setiawan, I. (2022) *Marketing 5.0: Technology for Humanity*. Hoboken, NJ: John Wiley & Sons.
- Lemon, K.N. and Verhoef, P.C. (2022) 'Understanding Customer Experience and the Customer Journey', *Journal of Marketing*, 86(6), pp. 69–96.
- Li, H. and Kannan, P.K. (2022) 'Attributing Conversions in a Multichannel Online Marketing Environment', *Journal of Marketing Research*, 59(3), pp. 461–479.
- Libai, B., Bart, Y., Gensler, S., Hofacker, C.F., Kaplan, A., Kohler, G. and Shankar, V. (2022) 'Brave New World? On AI and the Management of Customer Relationships', *Journal of Interactive Marketing*, 51(1), pp. 44–56.
- Matz, S.C., Teeny, J.D., Vaid, S.S., Peters, H., Harari, G.M. and Cerf, M. (2024) 'The potential of generative AI for personalized persuasion at scale', *Scientific Reports*, 14, 4692.
- McKinsey Global Institute (2023) *The Economic Potential of Generative AI: The Next Productivity Frontier*. New York: McKinsey & Company. Available at: <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai> (Accessed: 8 June 2024).
- Page, M.J., McKenzie, J.E., Bossuyt, P.M., Boutron, I., Hoffmann, T.C., Mulrow, C.D., Shamseer, L., Tetzlaff, J.M., Akl, E.A., Brennan, S.E., Chou, R., Glanville, J., Grimshaw, J.M., Hrobjartsson, A., Lalu, M.M., Li, T., Loder, E.W., Mayo-Wilson, E., McDonald, S., McGuinness, L.A., Stewart, L.A., Thomas, J., Tricco, A.C., Welch, V.A., Whiting, P. and Moher, D. (2021) 'The PRISMA 2020 statement: an updated guideline for reporting systematic reviews', *BMJ*, 372, n71.
- PwC (2023) *AI Predictions 2023: Now Decisions*. London: PricewaterhouseCoopers. Available at: <https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-predictions.html> (Accessed: 12 June 2024).
- Rust, R.T. and Huang, M.H. (2023) *The Feeling Economy: How Artificial Intelligence Is Creating the Era of Empathy*. Stanford, CA: Stanford Business Books.
- Salesforce (2023) *State of Marketing: 8th Edition*. San Francisco: Salesforce Research. Available at: <https://www.salesforce.com/resources/research-reports/state-of-marketing/> (Accessed: 10 June 2024).
- WEF (2023) *AI Governance Alliance: Briefing Paper Series*. Geneva: World Economic Forum. Available at: <https://www.weforum.org/publications/ai-governance-alliance-briefing-paper-series-2023/> (Accessed: 13 June 2024).