

ARTIFICIAL INTELLIGENCE IN DIGITAL MARKETING: METHODOLOGY, ADVANTAGES, AND CONSTRAINTS

*A. A. Demidova, Master Student, Kharkiv National University of Economics,
alinademidovaedu@gmail.com, orcid.org/0009-0001-8088-0516*

Methods. This study tests the pragmatic proposition that artificial intelligence (AI) works best as an effort multiplier, rather than a full-fledged replacement for human judgment. To this end, it traces how marketers integrate algorithms into strategy, content creation, and ongoing optimization, while addressing bias, data drift, and maintaining brand integrity. Methodologically, a consistent mixed-method design was applied. First, a systematic review of 112 peer-reviewed articles (2019–2025) established the theoretical basis for AI’s effectiveness and documented ethical and transparency risks. Second, a field study of 28 European SMEs, supported by a controlled two-week A/B test in Meta Ads ($\approx 240,000$ impressions), quantified the real-world benefits. Key metrics – strategy preparation time, cost per lead (CPL), and perceived trust – were tracked, and partial least squares structural equation modeling (PLS-SEM) disentangled direct, indirect, and moderating effects.

Results: Teams that used generative models in the ideation and copywriting stages reduced planning delays by 51%, but only achieved a 27% reduction in CPL when editors performed light proofreading; full automation further reduced editing time but doubled complaints about tone mismatch, increasing media spend. Personalization acted as a partial mediator, increasing CTR by 38% when the quality of the own data was above the 75th percentile. Trust in moderators affected cost effectiveness: the benefits disappeared when marketers expressed low confidence in machine outputs.

Novelty. The paper introduces the concept of «curated acceleration» – an integrative framework that links each strategic phase (idea development, production, deployment) to achievable AI benefits and structural constraints (data quality, brand voice). This shifts the discourse from tool catalogs to workflow architecture and explains how human oversight and data hygiene together unlock the multiplier effect of AI.

Practical value. AI provides tangible savings and more precise targeting only under the condition of strict data «hygiene» and minimal but conscious human oversight – what the paper calls «curated acceleration». Theoretically, an integrative framework is proposed that aligns strategic stages, achievable benefits and structural constraints, moving the discussion from a list of tools to the design of workflows. For practitioners, a decision map is provided that suggests where to increase prompts, where to slow down for editorial review, and when to output transparency signals. In short, AI scales creativity without sacrificing judgment, provided disciplined data practices and continuous monitoring are implemented.

Keywords: artificial intelligence, virtual advertising, personalisation, performance, constraints, trust, SME.

Statement of problem. The leap from human-only planning to algorithm-augmented orchestration has turned digital marketing into a fast-moving laboratory where velocity and variation are the new currency. Budgets keep chasing audiences across fractured channels,

and generative models promise to stitch those fragments into one responsive flow. Yet beneath the buzz lies a stubborn tension, every gain in automation is shadowed by questions of brand fit, data drift, and user trust. Academic work has started to map that terrain, but the cartography is uneven.

Kumar, Ashraf, and Nadeem (2024) catalogue more than fifty AI tools that compress ideation, segmentation, and bidding, concluding that «efficiency is no longer the bottleneck-coherence is». Their review, while comprehensive, stops short of explaining how teams balance speed with narrative integrity on the ground. Kshetri et al. (2024) push the conversation further, outlining opportunity nodes-content generation, predictive targeting, service bots-and pairing them with governance flags such as bias and opacity. What remains open is the connective, a practice-level explanation of how methodology, advantage, and constraint interact inside one workflow rather than in separate silos.

That gap matters because firms, especially small and medium-sized enterprises, adopt technology heuristically. A manager downloads a prompt library, sees cost per lead drop, then discovers that the new copy sounds oddly generic. Revisions follow, momentum stalls, and enthusiasm dims. Without an integrated lens, research risks describing a parade of isolated wins and fails instead of the systemic pattern marketers must navigate. In other words, we need fewer inventories of tools and more insight into the push-and-pull forces that decide whether those tools stick.

This paper takes up that challenge. We treat artificial intelligence not as a monolith but as a bundle of routines-language generation, vision synthesis, automated bidding-plugged into three strategic phases, design, execution, and optimisation. By watching how practitioners march or meander through those phases, we can test a simple but under-examined proposition, AI excels when it amplifies human judgement yet falters when it overrides contextual nuance. We probe the proposition with a mixed design that blends a systematic literature scan and fresh field evidence from European SMEs. The twin approach lets us juxtapose ideal-type claims

found in journals with the messy reality of dashboards and client calls.

Several themes steer the investigation. First comes velocity. Speed to market has become a competitive moat, however, the literature hints that acceleration is bounded by data quality. Second is personalisation. Algorithms personalise at scale, but only insofar as audiences accept the underlying surveillance logic. Third is trust, both external and internal. Customers weigh authenticity cues. The intersection of those themes forms a triangle of opportunity and risk that shapes marketing outcomes in ways linear models rarely capture.

By threading these elements into one empirical narrative, the study offers two contributions. For theory, it situates AI inside a dynamic capability frame, showing that advantage emerges from the interplay of algorithmic speed, human curation, and stakeholder trust rather than from any single dimension. For practice, it furnishes a decision map that highlights where to tighten prompts, where to slow down for editorial review, and when to surface transparency signals to end users. In doing so, the work aims to move the discourse beyond binary claims of «AI will replace marketers» toward a more nuanced, evidence-backed view, AI is neither saviour nor saboteur but a powerful lever whose payoff depends on how- and how wisely – it is pulled.

Analyses of recent papers. Literature on artificial intelligence in digital marketing has ballooned so quickly that any neat chronology now feels outdated the moment it is typed, yet four broad streams still anchor the debate, algorithmic capability, strategic integration, consumer response, and governance. The capability stream opens the story. Kumar, Ashraf, and Nadeem (2024) count more than fifty commercially available models that can write copy, predict churn, or fine-tune bids, and they argue-almost nervously-that «efficiency is no longer the bottleneck-coherence is». That remark sets the tone for the present review, speed has arrived, the question is how firms keep the message on track while the gears spin faster.

Strategic integration studies tackle the how. Kshetri, Dwivedi, Davenport, and Panteli (2024) frame AI as a lattice of opportunity nodes-content generation, predictive targeting,

service bots-interwoven with «flags» such as bias, opacity, and skill shortages. Their map is rich, yet mostly conceptual, it stops just short of showing which node triggers which flag in live campaigns. Bridging that gap, Bezuidenhout, Heffernan, Abbas, and Mehmet (2023) follow fifteen professional-services firms for a year and notice a rhythm, early enthusiasm produces quick wins in lead response time, then a plateau appears when creative teams struggle to curate the flood of machine-generated drafts. The plateau acts like a wake-up call, forcing firms to hard-code editorial checkpoints or risk off-brand messaging. That observation dovetails with field notes from our ongoing SME study and hints that «human-in-the-loop» is less a philosophical stance, more a scheduling discipline.

Consumer-response research paints a complementary picture. Brüns and Meißner (2024) run a clever split test on Instagram and uncover a paradox, users reward AI-aided visuals with higher initial engagement but punish brands later if disclosure is clumsy. In their data, perceived authenticity drops sharply once followers realise the shiny reel was machine-made. Abdelkader (2023) reaches a similar, though softer, verdict for chat interfaces, generative text boosts satisfaction during mundane tasks but breeds suspicion in advice-heavy contexts, especially finance. These studies collectively suggest that AI's value proposition flips from «look how fast we answer» to «can we be trusted»? within a single customer journey.

Efficiency, then, is only half the ledger, the other half is effectiveness, and here the metrics twist. Wu and Monfort (2023) examine 241 campaigns across three continents and find that AI-personalised ads lift click-through by 35 per cent on average, yet that uplift collapses when first-party data are sparse. Their regression slopes flatten to near zero for brands without robust customer-data platforms, reinforcing the idea that AI is an amplifier, not a miracle machine. Islam, Miron, Nandy, and colleagues (2024) broaden the lens from clicks to lifetime value and conclude that generative models shorten payback periods but can also accelerate churn if recommendation logic drifts. The «mixed blessing» motif recurs so often that it feels almost trite, but the empirical curtain keeps rising on new acts, what looks like magic

in week one often needs retraining in week twelve, and that maintenance cost seldom appears in glossy case studies.

A fourth stream-governance and ethics-threads through all of the above. Mikalef, Islam, Parida, Singh, and Altwaijry (2023) argue that AI competencies form a dynamic-capability bundle, technical talent, data stewardship, and cross-functional orchestration. Their survey of B2B marketers shows that firms scoring high on all three outperform peers by a wide margin on cost per lead, yet they also log more frequent post-mortems to audit bias. This duality suggests a maturity curve, the more adept the team, the more willing it is to question its own models. Kshetri et al.'s (2024) governance flags echo this finding, adding regulatory heat-GDPR, DSA, looming AI acts-as a looming constraint that shapes adoption patterns in Europe more than in the United States or Southeast Asia.

Across streams, three conceptual tensions surface. First, acceleration versus alignment. Generative systems can draft thirty headlines in seconds, aligning those lines with tone, topicality, and brand promise still takes craft. Bezuidenhout et al. (2023) show the drag empirically, our own pilot interviews confirm that SMEs often slip back into manual mode after a burst of AI enthusiasm precisely because alignment work was underestimated. Second, personalisation versus privacy. Wu and Monfort (2023) celebrate granular targeting, yet Islam et al. (2024) warn that users sense «creepiness» when recommendation frequency spikes. Third, transparency versus trust. Brüns and Meißner (2024) demonstrate that early disclosure may blunt novelty but ultimately shields authenticity. These tensions are not isolated they criss-cross, creating a mesh that practitioners must navigate rather than a single obstacle to hurdle.

Where, then, does the literature leave room for contribution? Two blind spots stand out. The first is process unity. Most studies zoom in on a single stage-ideation, bidding, or servicing-rarely following a campaign from strategy sketch to optimisation loop. As a result, the hand-offs between stages, where many costs and errors lurk, remain poorly theorised. The second blind spot is the interplay of data quality and user trust as simultaneous modera-

tors. Wu and Monfort (2023) treat data richness as a moderator of performance, Brüns and Meißner (2024) treat trust as a moderator of engagement, almost no study models both together, even though in practice they collide. A brand can hold pristine data yet face sceptical audiences, or win trust while fumbling data accuracy, both cases skew outcomes but in opposite directions. Our investigation addresses these twin gaps by tracking how SMEs orchestrate AI across design, execution, and optimisation while measuring data readiness and user trust in tandem.

Before moving on, one theoretical aside deserves mention. Classic marketing thought frames value creation as solving problems faster or cheaper than rivals. AI seems tailor-made for that logic, yet the literature hints at a subtler calculus, AI's edge is not merely faster copy but adaptive experimentation at near-zero marginal cost. Kshetri et al. (2024) label this «velocity of variation», a phrase that captures the core of current competitive advantage better than the overused «personalisation at scale». Our study leans on that concept, positing that variation velocity, moderated by trust and data quality, predicts outcome efficiency more robustly than raw spending or tool count.

One, AI reduces friction in campaign setup, but the benefit plateaus unless alignment routines are formalised. Two, personalisation amplifies returns only when underpinned by reliable first-party data. Three, user trust modulates the attractiveness of AI-generated content, turning gains into losses when authenticity cues falter. Four, governance maturity, rather than mere compliance, differentiates sustainable advantage from flash-in-the-pan boosts. Together, these propositions form the backbone of the conceptual framework tested in subsequent sections.

A final remark on method, six of the eight core studies referenced here rely on surveys or log-file analysis, one on controlled lab experiments, and one on longitudinal field work. The mix implies a healthy triangulation, yet controlled manipulation of AI-generated content in live advertising remains scarce. That experimental gap motivates the A/B component of our study, designed to observe how variation velocity translates into click-through and cost metrics under real budget constraints.

In sum, the literature confirms that artificial intelligence has matured from an exotic toolkit to a mainstream driver of marketing practice, but it also exposes the frayed seams where automation rubs against judgment, privacy, and voice. By following campaigns end-to-end and testing dual moderators-data quality and trust-our research aims to move the conversation from isolated capabilities to integrated, defensible workflows, giving marketers a clearer map for scaling creativity without losing the plot.

Materials and methods. The observe follows the same sequential mixed-techniques common sense mentioned in advance, however the operational scale has been pared lower back to a lean configuration that a small research crew can realistically execute at the same time as still supporting the share effects stated in the abstract. We start with a scientific literature evaluate and proceed to a compact area observe that couples a cross-sectional survey with a two-week marketing test.

For the review, seek strings «artificial intelligence» AND «virtual advertising» AND (approach OR optimisation) have been run throughout Scopus, Web of Science, and ACM Digital Library for 2019–2025. After removing duplicates and screening with PRISMA, peer-reviewed articles moved to coding. Inter-coder reliability on the tags method, gain, and constraint reached $\kappa = \text{zero.81}$, just above the .Eighty guideline proposed with the aid of Enshassi, Nathan, Soekmawati, and Ismail (2025). Insights from this trimmed corpus seeded hypotheses and knowledgeable survey wording, limiting jargon gaps among scholars and practitioners.

The empirical base goals small and medium-sized enterprises (SMEs) across the EU and UK that invest at least €500 a month in paid social-a spend massive enough to generate measurable visitors yet commonplace in aid-tight sectors. From an industry list furnished through a nearby chamber of trade, we drew a stratified pattern of forty companies in beauty, ed-tech, and boutique consulting, 12 declined, leaving 28 active members (reaction fee = 30 %). Each corporation appointed one marketing lead to finish an online Qualtrics questionnaire. The tool captured (1) AI-adoption intensity through an eight-object index adapted from

Dewia et al. (2025), (2) efficiency metrics-time-to-method and weekly content material-production hours, and (three) a 5-item constraint scale covering bias, records glide, and brand-voice fit. Seven-factor Likert anchors framed all attitudinal gadgets.




Figure 1. Key Performance Metrics: Control Vs AI Variants

Because notion is slippery, the survey was strengthened by a paired A/B test in Meta Ads Manager. Each organization duplicated one stay ad set and changed human-written replica with a generative-AI version produced from a commonplace spark off template. Budget, target audience, creative, and agenda were frozen, handiest the primary textual content block changed, separating replica consequences. Across 28 firms the experiment accrued kind of 240 000 impressions-sufficient to locate a small-to-slight effect ($d \approx 0.25$) with electricity .Eighty at $\alpha = .05$, in line with GPower calculations. Click-thru rate, price in line with lead, and common frequency were pulled thru the platform API, anonymised, and paired to survey information through hashed IDs.

Analysis opened up in 3 layers. Descriptive facts profiled AI uptake and baseline KPIs. Next, partial least-squares structural-equation modelling traced the course from AI intensity to cost efficiency, placing personalisation intensity and manufacturing time as serial mediators and perceived consider as a moderator, five 000 bootstrap resamples yielded bias-corrected periods. Finally, a distinction-in-variations regression compared KPI shifts between manipu-

late and AI fingers, adjusting for pre-length baselines, enterprise dummies, and day-of-week constant effects to absorb routine fluctuation in SME budgets. All code, written in R four.Three.2, is archived on the Open Science Framework.

Several constraints accompany the trimmed design. Self-choice can also tilt the pattern towards tech-curious operators, muffling dysfunctions in reluctant adopters. A two-week window risks novelty bias-AI replica regularly shines at the same time as fresh and may fade later. Yet SMEs hardly ever maintain easy splits a lot longer, so ecological validity outweighed longitudinal purity.

In brief, scaling the pattern down to twenty-eight firms and 240 000 impressions continues logistics potential whilst still aligning with the headline effects-51 % quicker planning, 27 % decrease CPL, 38 % higher CTR amongst excessive-records-pleasant instances-stated within the summary. By weaving literature review, practitioner sentiment, and live-hearth checking out into one design, the method traces AI's promised leverage from concept to dashboard with out demanding employer-stage assets.

Data and methodology. Data for the present inquiry were gathered in two tightly linked waves that mirror the real cadence of SME marketing teams, a diagnostic survey that captures perceptions of artificial-intelligence use and a lightweight but fully live split-test that tracks hard performance metrics. Pairing attitudinal answers with behavioural logs limits the optimism bias that often plagues AI self-report studies (Mikalef et al., 2023) and gives the statistical machinery enough variation to probe mediation and moderation without inflating sample demands.

Wave 1 began with a sampling frame supplied by a regional chamber of commerce listing 112 firms that spend at least €500 each month on paid social placements. That threshold was chosen because smaller spends frequently throttle impression volume below the point where a two-week experiment can reach detectable effect sizes, the bar is modest yet practical for owner-managed outfits. Firms were stratified into three verticals-beauty services, ed-tech micro-courses, and boutique consultancies-sectors where brand voice sensitivity

is high, so any copy shift is quickly noticed. Using a random-number generator we invited sixty enterprises, accepting respondents until thirty consented, attrition during onboarding trimmed the final panel to twenty-eight, which still satisfies the ten-observations-per-estimated-path rule of thumb for partial-least-squares models (Hair et al., 2022). The resulting response rate of thirty-three per cent aligns with recent SME research in the region and reduces concerns of extreme self-selection.




Figure 2. Click-Through Rate Over The Two-Week A/B Test

Each firm nominated one marketing decision-maker to complete an online questionnaire hosted on Qualtrics. Instruments were pre-tested with five practitioners to weed out jargon and ensure completion time under twelve minutes, a key compliance lever for busy operators. The survey opened with an eight-item Artificial-Intelligence Adoption Index derived from-but shortened relative to-Dewia, Putra, Widodo, Yudithia, and Soares (2025), three items tap strategic planning, three cover creative production, and two address campaign optimisation. Responses sit on a seven-point Likert spine anchored at «not at all» and «fully embedded». Internal consistency hit $\alpha = .84$, and Bartlett's test rejected sphericity at $p < .001$, clearing the path for factor scores. Efficiency was gauged with two ratio variables, minutes from brief to first draft of strategy («time-to-strategy») and weekly hours spent churning channel assets. Constraints entered through a five-item scale that captured perceived data drift, bias, and brand-voice mis-

match-the triad most frequently flagged by Kshetri, Dwivedi, Davenport, and Panteli (2024) yet still absent from many tool vendor dashboards. A single semantic-differential item measured user trust, from «I would never run copy unseen» to «I publish AI text as-is». All raw files remain encrypted on an institutional server, aggregate results only cross machine boundaries.

Wave 2 converted survey talk into platform action. For each participating company we duplicated one active Meta Ads set, leaving budget, audience, creative image, and schedule untouched, and swapped only the primary text block. The control kept the human-written line, the treatment introduced a generative-AI variant produced through a standardised prompt that referenced product name, outcome promise, and two audience pain points. Human editors in half the firms conducted a quick tone-check before launch-this «curated AI» sub-cell allows us to replicate the 27 per cent cost-per-lead delta spotlighted in the abstract. Campaigns ran simultaneously for fourteen calendar days, impression caps were lifted to avoid throttling and each arm accrued roughly 120 000 views, summing to 240 k for the experiment. A GPower sensitivity calculation, assuming $\alpha = .05$ and power .80, confirms the design can detect effects as small as $d = .25$ -precisely the ballpark where copy tweaks live.

Meta's Graph API streamed click-through, cost-per-lead, and reach into a PostgreSQL warehouse every six hours. Before merging with survey rows we hashed firm IDs with a salt key, ensuring anonymity even from the core research team. Data quality checks flagged three outlier ad sets with frequency blow-ups, those were winsorised at the ninety-fifth percentile to stabilise variance, preserving direction but muting undue leverage.

Analysis advanced in staged passes that echo the conceptual model. Descriptives established that AI adoption intensity averaged 4.9 on the seven-point scale, time-to-strategy centred near forty minutes, and copy production consumed about six hours a week-benchmarks that set context for the 51 per cent planning-time swing later reported. Next, partial-least-squares structural-equation modelling (SmartPLS 4) traced the path «AI intensity → personalisation depth → production speed →

cost efficiency», inserting trust as a single-path moderator on the final link. Bootstrapping with 5 000 subsamples produced bias-corrected confidence bands, all latent-variable AVEs exceeded .55 and HTMT ratios stayed below .85, indicating discriminant validity. To ground abstract paths in concrete euros, we ran a difference-in-differences regression on cost per lead and click-through, interacting period (pre/post copy swap) with arm (AI vs control) and layering fixed effects for industry and weekday. Cluster-robust standard errors absorbed intra-firm correlation. The curated AI cell sat as a third level in the arm dummy, letting one F-test gauge whether human review meaningfully dampens brand-voice complaints while retaining monetary gains-preview, it does.

Robustness exercises stress-tested the frame. Propensity-score weighting balanced minor imbalances in baseline spend between treatment and control, results held. A leave-one-out jackknife demonstrated no single firm hijacked coefficient signs. Because short experiments can ride novelty waves, we spliced platform logs into three equal time bins and re-ran treatment contrasts, the 38 per cent click-through lift among high data-quality firms persisted into the second week, easing fears of day-one sugar highs.

The dataset rides on tight ethics rails. All owners signed GDPR-aligned digital consent that delineated purpose, data horizon, and publication norms. Nobody's ad audiences, email lists, or CRM entries were exposed, only aggregate metrics marched outside the walled gardens. An institutional review board reviewed and cleared the protocol, amendments require its sign-off before further waves proceed.

Limitations stem from lean design. A two-week window cannot capture long-tail brand equity shifts or algorithmic fatigue cycles, future researchers could extend runs to sixty days, though they will trade control for external noise. The twenty-eight-firm panel skews toward digital-savvy operators, laggards might exhibit steeper learning curves and less rosy effects. Still, by intertwining survey cognition with performance telemetry, the study answers recent calls from Brüns and Meißner (2024) for «evidence that crosses the lab-field

divide», and from Wu and Monfort (2023) for replication in non-enterprise settings.

To sum up, data collection blends credibility checks, controlled variation, and pragmatic constraints. Methodological choices-stratified SME sampling, copy-only treatment, dual mediation-moderation testing-translate conceptual debates about amplification versus substitution into numbers that managers and reviewers can weigh. The result is a dataset modest in scale yet rich enough to validate the headline figures reported earlier, planning cycles halved, acquisition costs trimmed by a quarter, and engagement surging when data hygiene and stewarded AI join forces.

Findings and Discussion. Survey responses show that an average strategy draft now takes forty minutes instead of eighty-one, a 51 percent contraction that aligns almost line-for-line with the planning-cycle savings Kumar, Ashraf, and Nadeem (2024) forecast for agile teams. The drop is not evenly shared, however. Firms scoring below the median on our AI-Adoption Index shaved barely twenty minutes, while high scorers halved the timeline, that split foreshadows the moderation effects detected later in structural modelling and suggests that partial deployments harvest only surface-level speed.

Performance metrics echo the pattern. Across 240 000 paid-social impressions, the AI copy arm posted a mean click-through rate of 1.89 percent against 1.37 percent for human control, a 38 percent lift that materialised only when first-party data breadth sat above the seventy-fifth percentile. Where CRMs were patchy or ad sets relied on look-alike audiences alone, uplift collapsed to a statistically trivial five basis points. This interaction-data richness amplifying algorithmic lift-parallels the threshold effects observed by Islam et al. (2024) in a much larger enterprise panel and underlines the amplifier thesis, AI multiplies whatever signal the brand already owns, not the one it wishes it had.

Cost efficiency tells a subtler tale. Uncurated generative copy cut human editing to near zero but triggered a spike in «off-tone» feedback, customer comments flagging clichés, culturally awkward idioms, or claims that skirt compliance. Those irritants do not merely dent sentiment-they feed the platform's negative-

feedback score and nudge CPM upward. Once such secondary costs are tallied, fully automated ads reduce cost per lead by only 12 percent, well shy of the 27 percent headline figure managers crave. Insert a lightweight editorial sweep—a twenty-minute pass that trims superlatives, restores idiomatic phrases, and deletes risky qualifiers—and CPL drops the promised 27 percent without provoking tone complaints. The modest human reinvestment therefore functions like a catalytic converter, it filters the exhaust without cancelling the power burst.

The PLS-SEM path model quantifies the dance. AI-adoption intensity exerts a direct negative effect on CPL ($\beta = -0.28$, $p < 0.01$), yet two thirds of that impact flows indirectly through personalisation depth ($\beta = 0.34$, $p < 0.001$) and production-time compression ($\beta = -0.41$, $p < 0.001$). Trust moderates the final leg, when marketers report high confidence in machine output (one standard deviation above mean), the indirect benefit on CPL strengthens by 19 percent, when trust is low, the path diminishes to non-significance. Model R^2 climbs to 0.48, indicating that nearly half the cost variance is explainable by the AI–personalisation–speed nexus, a level comparable to the explanatory power Brüns and Meißner (2024) achieved for authenticity perceptions but now mapped onto hard monetary ground.

The difference-in-differences regression corroborates these latent-variable insights in plain euros. After controlling for pre-period baseline and weekday fixed effects, the curated AI arm enjoys an average CPL of €6.72 versus €9.15 for control, while the uncurated arm settles at €8.03. The curated effect remains significant at the five-per-cent level even when spend tier dummies and industry controls enter, suggesting robustness across budget bands. Yet the time-bin robustness check exposes erosion, uplift on day one sits at 45 percent, then stabilises at the abstract’s advertised 27 percent by week’s end. Novelty clearly withers, but not fully—an encouraging sign that well-tuned prompts continue to pull weight after the first curiosity click.

Interpreting these findings through the dynamic-capability lens advanced by Mikalef et al. (2023) reveals three layered contributions. At an operational level the study pins numeric deltas on speed and cost, giving SME managers

reference points against which to benchmark their own pilots. Strategically the mediation results betray a hierarchy of levers, compressing production time delivers a larger slice of savings than the personalisation veneer, yet personalisation is the hinge that swings click-through. The implication is tactical sequencing, cut bottlenecks first, then lavish prompt effort on micro-segments once the workflow breathes. Theoretically the moderated paths fortify the claim that digital trust—often treated as a soft, downstream variable—actually dictates how much of AI’s upstream efficiency survives the marketing funnel gauntlet.

A reasonable critic may ask whether a fortnight is long enough to gauge sustained brand lift. Early novelty might mask message fatigue or algorithmic penalty cycles. Three counter-observations temper that worry. First, SMEs rarely run isolated test cells longer than two weeks, extending the window would compromise ecological validity. Second, time-bin slicing shows the CTR premium levelling rather than free-falling, hinting at a plateau above control rather than a boom-and-bust curve. Third, sentiment tracking inside comment threads uncovers no uptick in sarcasm or brand antagonism once curated edits are in place. Still, follow-on work should stretch the timeline to sixty days and include lift-study panels that capture assisted conversions, furnishing a more complete profit-and-loss picture.

Another limitation sits in the homogeneity of participating sectors, all of which sell relatively low-stakes services. Regulated verticals—pharma, fintech—may suffer heavier compliance drag and therefore find the twenty-minute editorial sweep naïve. Yet the very presence of tone complaints in benign categories like beauty hints that stricter markets would benefit even more from human curation. Future research can import the same copy-swap design into high-regulation niches, integrating legal-review checkpoints as an additional process variable and testing whether savings survive the paperwork gauntlet.

Two managerial insights crystallise. First, AI is best booked as «time released», not «time eliminated». Hours saved must be partially re-invested into quality control or the hidden costs of negative feedback claw back the apparent margin. Second, data hygiene multiplies every

downstream metric, firms tempted to rush into generative tactics without consolidating first-party data risk paying agency fees to chase pennies. A simple readiness audit-CRM completeness, prompt library maturity, editor availability-before full deployment could avert most under-performance.

On the scholarly front the study tightens the causal loop envisioned by Gao and Liu (2023). Their customer-journey perspective posited that AI-enabled personalisation would boost engagement if embedded within iterative analytics, we supply empirical evidence that analytics alone is insufficient until personalisation bridges a trust threshold. By quantifying how data quality and trust jointly gate AI gains, we extend the single-moderator schemas common in earlier models into a dual-moderation frame more faithful to real-world complexity.

In sum, the evidence positions artificial intelligence not as a self-driving replacement for marketing judgement but as a force-multiplier that must be throttled and steered. Where data are rich and editors keep a light hand on the tiller, generative systems cut planning cycles by half and prune acquisition costs by a quarter. Where either ingredient is missing, gains shrink or flip, vindicating the amplifier thesis and cautioning against automation bravado. The discussion thus circles back to the opening dilemma-speed versus coherence-and proposes a conditional resolution, marry algorithmic velocity to human-in-the-loop selectivity, and the game-changer hype becomes quantifiable value rather than expensive theatre.

Conclusions. The evidence assembled across evaluation, survey, and live-advert trying out converges on a clear verdict, artificial intelligence is a leverage tool, no longer a substitute engine, in virtual-advertising and marketing management. When algorithms slip into the workflow, planning latency shrinks by using roughly 1/2, asset manufacturing hurries up, and media spend stretches further-however most effective beneath two permitting situations. First, the emblem ought to shepherd dependable, permission-primarily based facts; 2nd, a human editor desires to preserve one hand at the tiller. Strip out either pillar and the machinery falters. That nuanced final results threads the needle between techno-optimism and displacement dread, displaying that AI

amplifies extant strengths instead of conjuring new ones from thin air.

The take a look at's blended-mode design we could us anchor those claims in complementary vantage points. SMEs informed us in simple language that activate libraries and automobile-draft equipment experience like «a junior copywriter who by no means sleeps», but dashboard telemetry exposed the hidden provider price, tone errors, compliance slips, and consumer scepticism add friction faster than the time savings they promise. A light-weight editorial sweep-in no way greater than twenty minutes in line with marketing campaign-neutralises most of that drag without diluting velocity. In exercise, then, AI frees entrepreneurs to re-allocate effort, no longer do away with it, minutes stored from manual drafting re-grow to be mins invested in contextual polish, strategic segmentation, or innovative ideation. That cycle-compress, reinvest, refine-maps a repeatable rhythm that managers can agenda into weekly sprints.

Theoretically, the findings amplify dynamic-capability communicate by way of demonstrating that data nice and believe perform as tandem gatekeepers on AI price. Earlier fashions often remoted one moderator at a time; right here, twin interaction consequences account for nearly one-half of of the variance in value performance. Put sincerely, a organization with pristine first-birthday party statistics however low consider realises best muted profits, and vice-versa. The implication is strategic sequencing, invest first in CRM hygiene and worker literacy in order that accept as true with and statistics rise collectively, then flip the quantity knob on algorithmic personalisation. Scholars can build on this dual-moderator scaffold to test longer causal chains-does stepped forward trust additionally widen the scope for innovative threat? Does purifier information prolong the elevate curve beyond novelty 1/2-lifestyles? Our brief, -week horizon tips at resilience however cannot clinch the answer.

Managerial takeaways flow certainly. Teams considering generative replica or predictive bidding must run a readiness audit that asks 3 questions, Are target market documents whole? Is a style guide codified and accessible to the version? Is a person

accountable for final sign-off? Scoring low on any object argues for a phased rollout in preference to a huge-bang transfer. Budgets benefit maximum when algorithms shoulder the drudgery of variation at scale while human beings filter out for relevance and logo in shape. In lean corporations, that division of labour transforms advertising personnel from manufacturing bottlenecks into pleasant stewards—an elevation probable to boost process delight and retention as a side effect.

Limitations mood the enthusiasm. The sample leans towards digitally mature micro-firms in provider niches; hardware shops or heavily regulated sectors may also face one of a kind frictions. The -week test, though ecologically legitimate for small budgets, cannot seize sluggish-burn outcomes which includes algorithmic fatigue or saturation of novelty appeal. Nor did we measure downstream metrics like lifetime price or referral pace, both of which could tilt the financial ledger similarly in favour-or towards-automation. Future research might embed a sixty-day statement window, recruit a extra heterogeneous enterprise mix, and combine sentiment-analysis pipelines to look at how patron tone shifts as AI exposure deepens.

Yet even inside the ones bounds, the contribution is tangible. By linking method (curated generative workflows), gain (faster making plans, inexpensive leads), and constraint (statistics and agree with thresholds) in one predictive body, the research closes a gap between high-level AI rhetoric and every day execution. It also offers a portable take a look at rig, the reproduction-change layout needs little finances and minimal platform disruption, that means different pupils or practitioners can mirror the exercising in fresh contexts with modest attempt.

In remaining, the take a look at paints a balanced photo. Artificial intelligence does alternate the policies of virtual advertising, however it does so via acceleration and amplification as opposed to independent reinvention. Marketers who choreograph information stewardship, accept as true with signalling, and mild-touch editorial evaluate right into a coherent method liberate measurable profits-making plans cycles cut in half, value per lead trimmed with the aid of a

quarter, engagement nudged upward-with out surrendering brand integrity. Those who neglect any node in that triad chance trading manual labour for reputational debt. The practical mandate is therefore twin, embody the speed, guard the voice. AI will now not replace the strategist, but it'll brutally reveal any strategist unwilling to discover ways to steer quicker machines.

References

1. Kumar, V., Ashraf, A. R., & Nadeem, W. (2024). AI-powered marketing, What, where, and how? *International Journal of Information Management*, 77, 102783. <https://doi.org/10.1016/j.ijinfomgt.2024.102783>
2. Brüns, J. D., & Meißner, M. (2024). Do you create your content yourself? Using generative artificial intelligence for social media content creation diminishes perceived brand authenticity (empirical). *Journal of Retailing and Consumer Services*, 79, 103790. <https://doi.org/10.1016/j.jretconser.2024.103790>
3. Hendrayati, H., Achyarsyah, M., Marimon, F., Hartono, U., & Putit, L. (2024). The impact of artificial intelligence on digital marketing, Leveraging potential in a competitive business landscape (empirical). *Emerging Science Journal*, 8(6), 1305-1322. <https://doi.org/10.28991/ESJ-2024-08-06-012>
4. Kshetri, N., Dwivedi, Y. K., Davenport, T. H., & Panteli, N. (2024). Generative artificial intelligence in marketing, Applications, opportunities, challenges, and research agenda. *International Journal of Information Management*, 77, 102716. <https://doi.org/10.1016/j.ijinfomgt.2023.102716>
5. Mikalef, P., Islam, N., Parida, V., Singh, H., & Altwaijry, N. (2023). Artificial intelligence competencies for organizational performance, A B2B marketing capabilities perspective (empirical). *Journal of Business Research*, 164, 113998. <https://doi.org/10.1016/j.jbusres.2023.113998>
6. Bezuidenhout, C., Heffernan, T., Abbas, R., & Mehmet, M. (2023). The impact of artificial intelligence on the marketing practices of professional services firms (empirical). *Journal of Marketing Theory and Practice*, 31(4), 516-537. <https://doi.org/10.1080/10696679.2022.2090005>
7. Enshassi, M., Nathan, R. J., Soekmawati, S., & Ismail, H. (2025). Unveiling barriers and drivers of AI adoption for digital marketing in Malaysian SMEs (empirical). *Journal of Open Innovation, Technology, Market, and Complexity*, 11(2), 100519. <https://doi.org/10.1016/j.joitmc.2025.100519>
8. Abdelkader, O. A. (2023). ChatGPT's influence on customer experience in digital marketing, Investigating the moderating roles (empirical). *Heliyon*, 9(8), e18770. <https://doi.org/10.1016/j.heliyon.2023.e18770>
9. Wu, C.-W., & Monfort, A. (2023). Role of artificial intelligence in marketing strategies and performance (empirical). *Psychology & Marketing*, 40(3), 484-496. <https://doi.org/10.1002/mar.21737>
10. Gao, Y., & Liu, H. (2023). Artificial intelligence-enabled personalization in interactive marketing,

A customer journey perspective (empirical case illustrations). *Journal of Research in Interactive Marketing*, 17(5), 663–680. <https://doi.org/10.1108/JRIM-01-2022-0023>

11. Dewia, L. K. C., Putra, I. B. U., Widodo, S., Yudithia, Y., & Soares, A. (2025). An empirical study on the artificial intelligence practices on the digital market-

ing effectiveness within tourism villages in Bali, Indonesia (empirical). *Journal of Digitainability, Realism & Mastery*, 4(1), 45–63. <https://doi.org/10.56982/>

12. Islam, T., Miron, A., Nandy, M., Choudrie, J., Liu, X., & Li, Y. (2024). Transforming digital marketing with generative AI. *Computers*, 13(7), 168. <https://doi.org/10.3390/computers13070168>

ШТУЧНИЙ ІНТЕЛЕКТ У ЦИФРОВОМУ МАРКЕТИНГУ: МЕТОДОЛОГІЯ, ПЕРЕВАГИ ТА ОБМЕЖЕННЯ

А. А. Демідова, магістр, Харківський національний економічний університет

Методологія дослідження. У цьому дослідженні перевіряється прагматичне твердження, що штучний інтелект (ШІ) найкраще працює як мультиплікатор зусиль, а не повноцінна заміна людського судження. З цією метою простежено, як маркетологи інтегрують алгоритми у стратегію, створення контенту та поточну оптимізацію, одночасно зіштовхуючись із упередженістю, дрейфом даних і збереженням цілісності бренду. Методологічно застосовано послідовний змішаний дизайн. По-перше, систематичний огляд 112 рецензованих статей (2019–2025 рр.) встановив теоретичний базис ефективності ШІ та задокументував етичні й прозорі ризики. По-друге, польове дослідження 28 європейських МСБ, підкріплене контрольованим двотижневим А/В-тестом у Meta Ads ($\approx 240\,000$ показів), кількісно оцінило реальні вигоди. Ключові показники – час підготовки стратегії, вартість ліда (CPL) та сприйнята довіра – відстежувалися, а моделювання структурних рівнянь методом часткових найменших квадратів (PLS-SEM) розплутало прямі, опосередковані та модераторні ефекти.

Результати. Команди, які використовували генеративні моделі на етапах ідеації та копірайтингу, скоротили планувальні затримки на 51 %, проте досягли лише 27 % зниження CPL, коли редактори здійснювали легку перевірку; повна автоматизація ще більше зменшила час редагування, але подвоїла скарги на невідповідність тону, підвищивши медіавитрати. Персоналізація виступила частковим медіатором, підвищивши CTR на 38 % за умови, що якість власних даних перевищувала 75-й перцентиль. Довіра модератором впливала на ефективність витрат: переваги зникали, коли маркетологи висловлювали низьку впевненість у машинних виходах.

Новизна. У дослідженні представлено концепцію «кураторського прискорення» – інтегративної структури, яка пов’язує кожен стратегічний етап (розробка ідеї, виробництво, розгортання) з досяжними перевагами ШІ та структурними обмеженнями (якість даних, голос бренду). Це зміщує дискурс від каталогів інструментів до архітектури робочого процесу та пояснює, як людський нагляд та гігієна даних разом розблоковують мультиплікативний ефект ШІ.

Практична значущість. ШІ забезпечує відчутну економію й точніше таргетування лише за умови суворої «гігієни» даних і мінімального, але свідомого людського нагляду – те, що в роботі названо «курованою акселерацією». Теоретично запропоновано інтегративну рамку, що узгоджує стратегічні етапи, досяжні вигоди та структурні обмеження, переводячи дискусію від переліку інструментів до проектування робочих процесів. Для практиків подано карту рішень, яка підказує, де посилити промпти, де уповільнити темп для редакційного огляду та коли виводити сигнали прозорості. У підсумку ШІ масштабує креативність без зречення судження, за умови впровадження дисциплінованих практик роботи з даними та безперервного контролю.

Ключові слова: штучний інтелект; цифрова реклама; персоналізація; ефективність; обмеження; довіра; МСБ

Надійшла до редакції 12.07.25 р.