



# The Impact of Machine Learning Algorithms on the Effectiveness of Marketing Texts

Andrejs Zepkevicis

Founder & CEO, Epiic, Riga, Latvia.

## Abstract

*This article examines the impact of modern machine learning algorithms on the quality of marketing copy. Its importance lies in the growing role of text within digital marketing and the need for high-speed generation and testing of a large number of individualized messages. Falling cloud-computing costs and an open API ecosystem have made machine learning available to organizations of all sizes. At the same time, tightening regulations increase the demand for formal quality and security assessments of individualized content. This study shall perform a deep evaluation of generative transformers, segmentation tools (classification and clustering), and dynamic optimization algorithms (multi-armed bandits, Dynamic Creative Optimization) for the betterment of click-through rates and conversion rates via marketing communications. It will evaluate effectiveness at every stage, from text variant generation through automatic live testing to real-time traffic reallocation. The novelty of this study lies in its integration of three beneficial technological components into a single looping design: generative tools that provide text variety, split methods that sharply focus on key readers, and rapid fix systems that modify content instantly using user input. Ideas are presented alongside real-world outcomes from more than 16 diverse sources, including examples from JPMorgan Chase, Pizza Hut, Netflix, and Nespresso. Key findings indicate that using transformers and DCO can give a 4.5 times increase in click-through rate and an 8-12% rise in conversion. In return, contextual bandits can lead to a growth of up to 30% in deals without long A/B test cycles. At once, four major dangers are found: model illusions, algorithmic prejudice, data-privacy worries, and content 'out-of-dateness' because of overuse, each needing truth-check layers, fairness reviews, strong privacy safeguards, and ways to bring new creative ideas. The article will be useful for marketing professionals, data analysts, and AI solution developers.*

**Keywords:** Machine Learning; Marketing Texts; Generative Transformers; Segmentation; Multi-Armed Bandits; Dynamic Creative Optimization; Personalization; ROI; A/B Testing.

## INTRODUCTION

Text remains the primary interface between a brand and an individual: it is at the level of words that the first impression of a product is formed, the cognitive evaluation of its value is initiated, and—when motivation and context align—a conversion occurs. Even in visual channels such as Instagram or TikTok, a follower reads the video headline and the lexical form of the call-to-action before tapping the button. Marketing research confirms this text-centric core of the funnel: a personalized email subject line increases the probability of open rates by 26%, with the effect observed across large samples from both e-commerce and B2B companies [1]. Thus, any improvement in text quality or relevance immediately impacts top-of-funnel metrics and, through cascading influence, the overall campaign ROI.

The need to craft ever more precise and varied messages confronts human limitations: a creative team might propose

dozens of formulations, whereas the digital environment demands thousands. Consequently, the current cycle of AI-driven content transformation began not with analytics but with generative models. A recent Salesforce survey reveals that 51% of marketers have already adopted or are actively testing generative AI, and 22% plan to implement it within the next year. Of those currently using it, 76% apply models to tasks such as writing ad copy and basic content creation [2]. The market has thus embraced AI as a regular part of the creative process: algorithms take on the grunt work and draft ideas; humans concentrate on final stylistic polish and semantic checks.

The second technological lever is the convergence of machine learning with dynamic creative optimization. Market growth necessitates that automatic testing of text variants in live conditions becomes not an exotic experiment but a prerequisite for competitiveness: algorithms can redirect

**Citation:** Andrejs Zepkevicis, "The Impact of Machine Learning Algorithms on the Effectiveness of Marketing Texts", Universal Library of Business and Economics, 2025; 2(3): 47-51. DOI: <https://doi.org/10.70315/uloap.ulbec.2025.0203008>.

traffic to the winning formulation within hours, whereas classical A/B testing requires days or even weeks.

Ultimately, the macroeconomic context necessitates that companies accelerate their adoption. Declining cloud-computing costs and an open API ecosystem render model access nearly barrier-free; concurrently, regulatory pressures on personalization quality and data protection intensify, driving the shift from heuristics to formal ML approaches. As a result, sales-funnel text becomes a field of rapid evolution: every word can be generated, evaluated, and rewritten by machines in real time, and campaign effectiveness is measured not by the number of characters produced but by statistically significant gains in business outcomes.

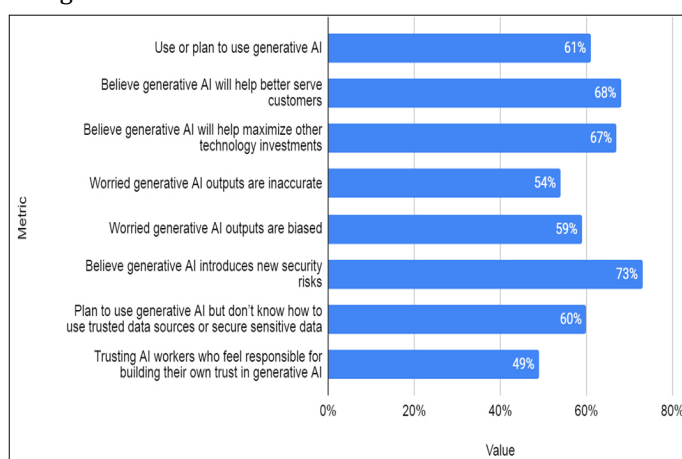
## MATERIALS AND METHODOLOGY

The materials and methodology of this research are grounded in a comprehensive analysis of 16 sources, including academic articles, industry reports, case studies, and marketer surveys. The theoretical foundation comprises publications on text personalization and AI content generation: an American Marketing Association study confirmed the influence of personalized subject lines on email open rates [1], and a Salesforce report identified the proportion of marketers already using or testing generative AI in copywriting and creative tasks [2]. Data on the dynamic growth of the Dynamic Creative Optimization (DCO) market were derived from a Business Research Insights report [3]. At the same time, practical outcomes of early implementations were drawn from a Brandxr review of generative AI applications in marketing [4] and from e-commerce experiments by H. Yang et al. [5]. The concept of a fast feedback loop via multi-armed bandit algorithms and DCO is supplemented by McKinsey's findings on ROI from text automation [6] and by Grohaus' case studies on deploying bandit experiments in lifecycle marketing [9].

The methodological approach comprised several stages. First, a comparative text-generation model evaluation of the gap between transformers and old template systems, and how well fine-tuning or RLHF works on proprietary data. Second, an empirical analysis of A/B tests and bandit experiments showing CTR and conversion gains in the JPMorgan Chase case with Persado [4], Pizza Hut via Braze [14], and Netflix with a contextual bandit [10]. Third, a systematic review of the DCO market, including an assessment of CAGR through 2033 [3], was combined with marketer surveys by Salesforce [2] and SurveyMonkey [13], enabling the alignment of stated implementation plans with actual performance metrics. Finally, a content and case analysis of Dynamic Creative Optimization platforms—exemplified by Nespresso [16]—was conducted to identify mechanisms for real-time creative assembly and optimization. This multi-layered approach provided a holistic understanding of how machine learning algorithms influence the effectiveness of marketing texts, allowing for the juxtaposition of theoretical expectations with practical business results.

## RESULTS AND DISCUSSION

The algorithmic core of the modern content engine is built around two complementary classes of models, each addressing precisely those narrow tasks identified in the Introduction as the volume and relevance bottleneck of texts. Generative transformers, trained on trillions of tokens, create a virtually unlimited number of different ways to say something. They don't rely on scripts or spinning of words as was done before. Rather, they generate new linguistic patterns through self-attention mechanisms and large-scale, pre-trained contextual representations; later adaptation occurs either by fine-tuning on proprietary datasets or via RLHF, which explores real user behavior. A Salesforce survey shows the scale of adoption: 61% of knowledge workers are already using generative AI in their work-related communications or plan to do so in the near future, and 68% believe these tools will help them better serve their customers [2], as illustrated in Fig. 1.



**Fig. 1.** Examining Employee Adoption, Perceived Benefits, and Trust Constructs in Generative Artificial Intelligence Utilization [2]


The outcome goes beyond what one would anticipate subjectively. JPMorgan Chase, collaborating with Persado, attained a 4.5x increase in CTR over human copywriting, with preliminary results obtained in less than a quarter [4]. Similar trends were unveiled in academic validation: an experimental LLM-generation framework for e-commerce led to a 12.5% increase in clicks and an 8.3% increase in conversions without compromising textual originality [5]. McKinsey extrapolates that the cumulative revenue uplift from such text enhancements and related automation can deliver a business ROI of 10–20% above baseline sales trajectories [6]. In practice, these models simultaneously address three tasks: the creation of automatic draft copies, the generation of micro-headlines for different traffic segments, and the rapid rewriting of messages based on online test outcomes.

However, even the most persuasive text will fail if shown to the wrong audience, so the second algorithmic block is devoted to precise targeting. Classical classification methods (e.g., gradient boosting over RFM features) generate a

response-probability forecast and automatically incorporate these scores into display rules. More fundamentally, clustering algorithms—where segments emerge from the data rather than fixed labels—solve the problem via K-means, Gaussian Mixture Models, and, in complex cases, spectral and RL-reinforced variants. The observed effect rivals that of content generation but is achieved by reducing noisy exposures: a banking study showed that moving from demographic segmentation to K-means clustering produced a marked increase in engagement and positively correlated with campaign profitability [7], and in a retail example, ZALORA's implementation of ML-driven micro-segments via Twilio Segment doubled the conversion rate by delivering more relevant triggers to shoppers [8]. In combination, classification solves the to whom filter, while clustering answers the question of which motive, and together with transformers, they form a closed loop: the model generates text for a specific motive, and the segmentation system delivers it to those for whom that motive is statistically significant.

The following logical link—after text generation and precise segmentation—is their instantaneous adaptation to user behavior, which is provided by a pair of contextual bandit algorithms and reinforcement-learning methods. The model treats each message variant as an arm of a one-armed bandit and reallocates traffic in real-time to formulations yielding higher immediate responses, while preserving exploration traffic. In corporate practice, this has moved beyond laboratories: Pizza Hut, running bandit experiments on the Braze platform, increased transactions by 30%, revenue by 21%, and profit by 10% without a separate classical A/B-testing cycle [9]. Netflix-scale systems handle up to twenty million personalized image requests per second, treating this stream as a contextual-bandit problem where the reward is the probability of starting to watch the selected title [10].

Dynamic Creative Optimization (DCO) extends this concept to the entire ad unit, combining text, visuals, and pricing in hundreds of milliseconds. The DCO market was already valued at \$0.9 billion in 2024 and is forecasted to grow to \$2.19 billion by 2033 at a 10.2% CAGR [3], as shown in Fig. 2.

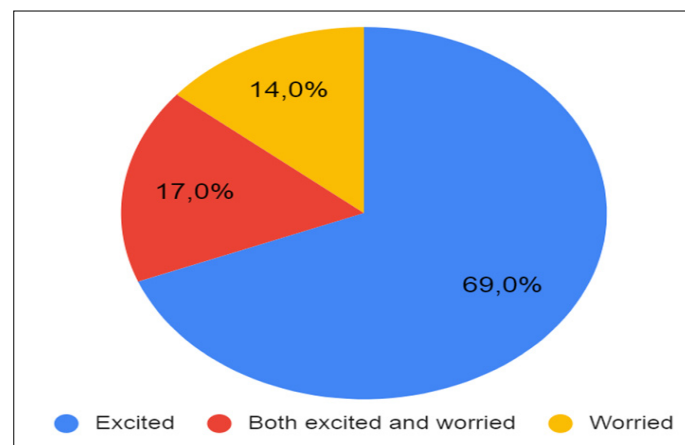


**Fig. 2.** Projected Market Valuation Dynamics Under a 10.2% Compound Annual Growth Rate [3]

Practice confirms the commercial payoff: in native advertising on Yahoo Gemini, the introduction of an explore-exploit mechanism for selecting combinations of ad elements resulted in a 53.5% uplift in the conversion rate relative to uniform variant allocation [11]. Surveys also document the mass adoption of the technology: 82% of brands already employ DCO, one-third plan to increase their investment, and a further 49% intend to maintain at least current levels—evidence of the method's firm entrenchment in media plans for the coming years [12].

Thus, contextual bandit algorithms and DCO establish a fast loop around the previously described generation-and-segmentation pipeline: a transformer generates numerous hypotheses, clustering assigns them to target recipients, and a bandit or reinforcement-learning agent instantly reallocates impressions to maximize click or purchase likelihood; DCO subsequently extends this principle to the multimodal components of an ad. The outcome is a sustained rise in key KPIs without the delays inherent in classical A/B testing cycles, laying the groundwork for fully autonomous marketing systems in which each new message immediately contributes to the model's continuous learning.

The shift from static creative development to algorithmically driven communications has transformed the nature of marketing effectiveness: now, each word lives only as long as it is read. Mass personalization has been the earliest visible result. Given that even the simple insertion of a recipient's name into an email subject line raises open rates by 26%—as noted in the introduction—it is unsurprising that 73% of marketers already leverage AI for constructing personalized experiences, and 88% do so as part of their daily content workflows [13]. Generative models supply a steady stream of creative variants, segmentation algorithms distribute them across micro-audiences, and automated delivery orchestrates real-time hypothesis testing. Most respondents report a generally positive perception of AI's impact on their work, whereas only 14% feel primarily anxious and 17% describe a mixed state of excitement and anxiety.



**Fig. 3.** Distribution of Emotional Responses to AI Integration and Its Perceived Impact on Occupational Roles [13]

Automated message-market fit is achieved through continuous formulation testing. Contextual multi-armed

bandit algorithms deployed at Pizza Hut reallocated traffic among hundreds of variants and, within days, identified top performers—yielding a 30% increase in transactions, 21% in revenue, and 10% in profit without traditional split-testing [14]. This approach marries exploitation of high-performing copy with controlled exploration of new alternatives, securing steady metric growth while eliminating the manual write–wait–analyze cycle.

Resource savings emerged as the second most significant benefit. Fintech firm Klarna demonstrated that generative AI can both lower creative costs and compress production timelines. Shifting to automated image and copy generation reduced external marketing expenditures by 11%, achieved annual savings of approximately \$10 million, and shortened visual preparation from six weeks to seven days [15]. Budget-wise, this enables more campaign launches at constant cost; team-wise, it liberates creative capacity for strategic initiatives.

The contraction of the idea → live campaign cycle reaches its zenith in Dynamic Creative Optimization systems. In the Nespresso case, a DCO platform assembled text, price, and imagery in real-time, driving a 38% uplift in coffee-machine sales, a 15% increase in capsule sales, and concurrently reducing customer-acquisition costs by 36% [16]. Here, the generative module ensures variant diversity, the bandit core executes rapid winner selection, and the DCO layer instantaneously composes channel- and user-specific creatives. The result is not merely enhanced performance metrics but a paradigm shift to marketing as an uninterrupted model-training process, where each click refines the precision of subsequent messages.

Machine learning enhances marketing performance; however, its very characteristics give rise to four systemic limitations that become apparent immediately upon model deployment and directly affect trust in automated copywriting.

First is the risk of hallucinations: language models, even the most advanced, sometimes lack factual information and provide incorrect details about a brand. It leads to inaccuracies that can shatter reputational capital in the advertising context; hence, manual validation costs negate any speed advantage. For this practical reason, companies must add an automatic fact-checker layer and codify a refusal procedure for the model when the veracity of a claim cannot be determined.

The second limitation concerns ethical biases. Algorithms trained on historical advertisements and user dialogues inevitably inherit the same gender, cultural, and social stereotypes. As a result, an identical message can provoke markedly different reactions across audience segments, and without monitoring fairness metrics, the risk of reputational damage increases. This issue cannot be resolved by merely cleansing the dataset; it requires regular model audits that include a set of benchmark cases and specialized fairness metrics.

The third risk pertains to data confidentiality. Text personalization relies on quite sensitive information—behavior and purchase history. A leak of this data not only threatens fines but, more importantly, erodes user trust. Thus, companies are slowly moving part of the inference process to client devices, adding differential privacy and encrypting data at the transport layer. However, while this enhances security, it complicates DevOps workflows, as the model must respond within milliseconds.

Finally, the mechanics of continuous optimization inherently lead to formulaic output: when a contextual bandit persistently reinforces the most clickable variant, the system gradually suppresses unconventional formulations. After several weeks, a campaign may degenerate into a repetitive set of phrases to which the audience becomes desensitized, and conversion rates begin to decline. Mitigating this drift requires the periodic introduction of fresh hypotheses and a creative buffer—a portion of traffic reserved for experimental variants, even if they temporarily underperform on established metrics.

Thus, maximizing effectiveness requires not only accurate models but also a comprehensive risk-mitigation architecture: veracity filters, fairness audits, robust data-protection schemes, and mechanisms for continual creative renewal. Only by simultaneously controlling all four risk domains does growth in clicks and conversions truly translate into a sustainable competitive advantage for the brand.

## CONCLUSION

The present study demonstrates that modern machine-learning algorithms fundamentally transform the process of creating and optimizing marketing copy, elevating it to a qualitatively new level of efficiency. Generative transformers provide a continuous stream of variant formulations adaptable to brand specifics and audience behavior, while classification and clustering methods enable precise message delivery to relevant segments. Contextual multi-armed bandit algorithms and Dynamic Creative Optimization (DCO) systems, in turn, implement a fast feedback loop, instantly reallocating traffic in favor of the most successful variants and thereby eliminating the delays of classical A/B testing.

Practical case studies confirm the significance of this approach: JPMorgan Chase's partnership with Persado resulted in a 4.5-fold increase in CTR compared to traditional copy, and e-commerce experiments demonstrated a 12.5% uplift in click-through rates and an 8.3% conversion rate. Pizza Hut's deployment of contextual bandits increased transactions by 30%, revenue by 21%, and profit by 10% without the time investment of a separate split test.

Thus, making the best use of machine-learning-based marketing writing comes from a joint design where creation, division, and real-time check make unbreakable parts of a single loop. The effective use of such setups requires not only precise models and substantial computer power but also comprehensive risk management, which encompasses truth-



check tools, fairness reviews, data security, and maintaining creative diversity. Only with this balanced strategy can copywriting automation deliver a sustainable competitive advantage and long-term growth in the key metrics of marketing campaigns.

### REFERENCES

1. American Marketing Association, "An Expert Guide to Email Personalization," American Marketing Association, May 01, 2024. <https://www.ama.org/marketing-news/email-personalization-strategies/> (accessed May 21, 2025).
2. "Top Generative AI Statistics for 2025," Salesforce, 2025. <https://www.salesforce.com/news/stories/generative-ai-statistics/> (accessed May 22, 2025).
3. "Dynamic Creative Optimization Market Size," Business Research Insights, 2025. <https://www.businessresearchinsights.com/market-reports/dynamic-creative-optimization-dco-market-112176> (accessed May 23, 2025).
4. M. Mattan, "Generative AI in Marketing: Transforming Content Creation," Brandxr, 2025. <https://www.brandxr.io/the-complete-ai-marketing-playbook-from-strategy-to-implementation> (accessed May 24, 2025).
5. H. Yang, H. Lyu, T. Zhang, D. Wang, and Y. Zhao, "LLM-Driven E-Commerce Marketing Content Optimization: Balancing Creativity and Conversion," arXiv, 2025. <https://arxiv.org/abs/2505.23809> (accessed May 25, 2025).
6. R. Deveau, "Marketing and Sales Soar with Generative AI," McKinsey, May 11, 2023. <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/ai-powered-marketing-and-sales-reach-new-heights-with-generative-ai> (accessed May 26, 2025).
7. "Customer Segmentation Using K-Means Clustering for Personalized Marketing Campaigns," International Journal of Scientific Research in Computer Science, Engineering and Information Technology, vol. 11, no. 3, pp. 810–815, 2025, doi: <https://doi.org/10.32628/cseit25113344>.
8. "How ZALORA increased experimentation velocity and doubled its conversion rates with Twilio Segment," Twilio Customer Stories, 2025. <https://customers.twilio.com/en-us/zalora> (accessed May 29, 2025).
9. "The Commoditization of Multi-Armed Bandit Experiments for Lifecycle Marketing," Grohaus, 2024. <https://www.grohaus.io/content/multi-armed-bandit-experiments-for-lifecycle-marketing> (accessed May 29, 2025).
10. F. Amat, A. Chandrashekar, T. Jebara, and J. Basilico, "Artwork personalization at Netflix," Proceedings of the 12th ACM Conference on Recommender Systems - RecSys '18, 2018, doi: <https://doi.org/10.1145/3240323.3241729>.
11. Y. Kaplan, Y. Koren, A. Shtoff, T. Shadi, and O. Somekh, "Conversion-Based Dynamic-Creative-Optimization in Native Advertising," arXiv, 2022. <https://arxiv.org/abs/2211.11524> (accessed Jun. 01, 2025).
12. C. Mann, "Survey: DCO spend surge predicted," Advanced Television, Mar. 07, 2024. <https://www.advanced-television.com/2024/03/07/study-dynamic-creative-optimisation-spend-surge-predicted/> (accessed Jun. 02, 2025).
13. "28 AI marketing statistics you need to know in 2025," SurveyMonkey, 2024. <https://www.surveymonkey.com/mp/ai-marketing-statistics/> (accessed Jun. 03, 2025).
14. "Pizza Hut Adds Key New Channels and Drives a 21% Increase in Revenue With Machine Learning," Braze, 2025. <https://www.braze.com/customers/pizza-hut-case-study> (accessed Jun. 04, 2025).
15. C. Kelly, "Klarna uses AI to save \$10M on marketing annually while upping output," Marketing Dive, May 29, 2024. <https://www.marketingdive.com/news/klarna-gen-ai-openai-cut-marketing-spend-efficiency/717332/> (accessed Jun. 06, 2025).
16. "Segmenting Audiences and Increasing Sales with A/B Testing at Scale," Clinch, 2025. <https://www.clinch.co/case-studies/nespresso> (accessed Jun. 07, 2025).