

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

International Journal of Research in Marketing

journal homepage: www.elsevier.com/locate/ijresmar

Full length article

The power of generative marketing: Can generative AI create superhuman visual marketing content?

Jochen Hartmann ^{a,*}, Yannick Exner ^{a,1}, Samuel Domdey ^b^a Technical University of Munich, TUM School of Management, Arcisstr. 21, 80333 Munich, Germany^b Technical University of Berlin, Straße des 17. Juni 135, 10623 Berlin, Germany

ARTICLE INFO

Article history:

First received on 23 July 2023 and was under review for 8½ months

Available online 20 September 2024

Area Editor: Lan Luo

Accepting Editor: David A. Schweidel

Keywords:

generative AI
marketing effectiveness
productivity
content creation
artificial intelligence
digital marketing

ABSTRACT

Generative AI's capacity to create photorealistic images has the potential to augment human creativity and disrupt the economics of visual marketing content production. This research systematically compares the performance of AI-generated to human-made marketing images across important marketing dimensions. First, we prompt seven state-of-the-art generative text-to-image models (DALL-E 3, Midjourney v6, Firefly 2, Imagen 2, Imagine, Stable Diffusion XL Turbo, and Realistic Vision) to create 10,320 synthetic marketing images, using 2,400 real-world, human-made images as input. 254,400 human evaluations of these images show that AI-generated marketing imagery can surpass human-made images in quality, realism, and aesthetics. Second, we give identical creative briefings to commissioned human freelancers and the AI models, showing that the best synthetic images also excel in ad creativity, ad attitudes, and prompt following. Third, a field study with more than 173,000 impressions demonstrates that AI-generated banner ads can compete with professional human-made stock photography, achieving an up to 50% higher click-through rate than a human-made image. Collectively, our findings suggest that the paradigm shift brought about by generative AI can help advertisers produce marketing content not only faster and orders of magnitude cheaper but also at superhuman effectiveness levels with important implications for firms, consumers, and policymakers. To facilitate future research on AI-generated marketing imagery, we release [GenImageNet](#) that contains all of our synthetic images and their human ratings.

© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Generative AI fundamentally disrupts the marketing industry, representing a new paradigm of automated marketing content generation (Peres, Schreier, Schweidel, & Sorescu, 2023). Industry reports suggest a tremendous economic potential of generative AI, quantifying its impact at USD 463 billion in the marketing sector alone (Chui et al., 2023). Both marketing practice and research report astonishing anecdotal examples of generative AI's disruptive possibilities (Kelly, 2023; Noy & Zhang, 2023). Encouraged by such promising prospects, some firms have already successfully piloted synthetic content created by generative AI in their marketing campaigns (Acar & Gvirtz, 2024), e.g., the award-winning "A.I. Ketchup" campaign by Heinz, which garnered more than 850 million earned impressions around the globe (The One Club, 2023).

* Corresponding author.

E-mail address: jochen.hartmann@tum.de (J. Hartmann).¹ These authors contributed equally to this work and share first authorship.

Given the considerable excitement around generative AI, it is not surprising that firms have started exploring and experimenting with this novel technology. Industry forecasts project that large organizations will synthetically generate up to a third of their outbound marketing messages by 2025 (Gartner, 2024). However, the sustainable adoption of generative AI by firms critically hinges on generative AI's effectiveness in reaching their marketing objectives (Jansen, Heitmann, Reisenbichler, & Schweidel, 2024) and its efficiency, namely, in realizing substantial cost savings (Ammanath, Dutt, Perricos, & Sniderman, 2024; Gartner, 2024). Pioneering studies demonstrate the productivity gains and increase in output quality enabled by generative AI for automated marketing text generation (e.g., Reisenbichler, Reutterer, Schweidel, & Dan, 2022; Reisenbichler, Reutterer, & Schweidel, 2023). Preeminent studies outside of marketing corroborate these generative AI-enabled improvements with tangible economic benefits (e.g., Noy & Zhang, 2023; Brynjolfsson, Li, & Raymond, 2023). However, due to the recency of the "age of generative AI" (Krugmann & Hartmann, 2024) and idiosyncratic challenges pertaining to image creation (Borji, 2023), little is known about its disruptive potential for visual marketing content across diverse marketing contexts.

A better understanding of AI-generated marketing imagery's effectiveness and efficiency is important as images are a cornerstone of today's marketing communications in an increasingly media-rich environment (Grewal, Gupta, & Hamilton, 2021). Firms and their ad agencies carefully design online and offline ads (Pieters & Wedel, 2004; Hartmann, Heitmann, Schamp, & Netzer, 2021), influencers get paid to endorse brands across visual social media channels (Beichert, Bayerl, Goldenberg, & Lanz, 2024), online shops present products and services in the best possible conditions (Dzyabura, El Kihal, Hauser, & Ibragimov, 2023; Zhang, Lee, Singh, & Srinivasan, 2022b), consumers share their everyday consumption experiences online (Li & Xie, 2020; Zhang & Luo, 2023), and their digital traces offer a wealth of information for brand managers to visually "listen in" (Liu, Dzyabura, & Mizik, 2020; Dzyabura & Peres, 2021). How do consumers perceive and respond to synthetic images compared to human-made content? How does AI-generated marketing imagery perform in a real-world context? If generative AI could create human-level visual marketing content, it could fundamentally challenge traditional human-made marketing content generation and accelerate AI adoption.

The importance of generative AI's role for the future of marketing is underscored by the substantial cost associated with creating professional marketing imagery, especially when considering large-scale, global marketing campaigns, which can require hundreds of visual assets tailored to different communication channels and target audiences (King, 2024). Take the following examples: Purchasing a professional stock photo typically costs around USD 5 – 10, excluding additional expenses to acquire more permissive usage licenses. Opting for an experienced freelancer from an online marketplace to create a custom marketing image can increase the cost by an order of magnitude to around USD 100. Employing top-tier ad agencies or organizing professional photo shoots, which involves specialized photographers and cast photo models, can even result in expenses ranging from thousands to tens of thousands of USD (Rodgers, 2021). In contrast, generating a single image with OpenAI's DALL-E 3, a state-of-the-art generative text-to-image model, costs merely USD .04 (Betker et al., 2023).²

What if generative AI could substantially lower the expenses associated with the time-consuming and cost-intensive process of creating marketing imagery without compromising the content's visual appeal and marketing effectiveness? Is a prompt consisting of a couple of words and the right AI model³ all an advertiser needs? Considering that most methods are developed in computer science as general-purpose AI tools without specific optimization for marketing applications (Bommasani et al., 2021; Dzyabura, El Kihal, & Peres, 2022), it is unclear if state-of-the-art generative text-to-image models can generate effective marketing content that resonates with consumers when used off the shelf. Similarly, there is a lack of scientific evidence on which AI models provide consistent performance across marketing applications.

To systematically address this research gap, we conduct three studies. First, we investigate the perceptual evaluation of AI-generated vs. human-made marketing images. Study 1 draws on eight different real-world marketing datasets, covering a comprehensive set of marketing applications, structured by the source of the data (firms vs. users) and the marketing objective (call to action vs. convey brand identity). We prompt seven state-of-the-art generative text-to-image models, released between October 10, 2023 and February 1, 2024, namely, DALL-E 3, Midjourney v6, Firefly 2, Imagen 2, Imagine, Stable Diffusion (SD) XL Turbo, and Realistic Vision to generate 10,320 synthetic images, using 2,400 real-world, human-made images as input. 254,400 human evaluations of these images, combined with algorithmic aesthetics assessments, show that AI-generated marketing imagery can surpass human-made images in quality, realism, and aesthetics.

Second, we give identical creative briefings to commissioned human freelancers and the same AI models, mimicking a real-world advertising pretest (MacKenzie, Lutz, & Belch, 1986). We evaluate the perception, attitude, behavioral intention, and prompt following of the AI models and human freelancers in a between-subjects design across ten dependent variables (N = 1,575 Prolific panelists). Overall, DALL-E 3 produces the best synthetic images, outperforming the human freelancers in terms of five marketing metrics, and obtaining directionally higher evaluations across the other five. Strikingly, participants attribute higher ad creativity to the AI-generated images by DALL-E 3 compared to the human-made freelancer images. In addition, AI-generated images are substantially more cost-efficient. The same budget of a single freelancer image allows for creating 2,500 images with DALL-E 3.

² See <https://openai.com/pricing> (accessed September 2, 2024)

³ We use the terms "AI model" and "generative text-to-image model" interchangeably to refer to systems that make use of AI to generate images based on textual descriptions. Similarly, "AI-generated" and "synthetic" images are used synonymously (see also <https://deepmind.google/discover/blog/identifying-ai-generated-images-with-synthid/>; accessed April 4, 2024).

Third, we run a real-world marketing campaign on an online marketing platform to analyze the actual effectiveness of AI-generated banner ads in terms of their click-through rates (CTRs). We collect and evaluate over 173,000 impressions to compare the synthetic images with a high-quality, human-made stock photo selected by an online marketing professional. DALL-E 3, the best-performing AI model, achieves an over 50% higher CTR than the human-made banner ad, while being 225 times cheaper to create. DALL-E 3's CTR significantly outperforms the least effective AI model, SDXL Turbo, by 100%.

This research makes three important contributions. First, we shed light on the real-world marketing effectiveness of AI-generated vs. human-made images. While nascent marketing research demonstrates generative AI's effectiveness for textual content generation (e.g., Carlson, Kopalle, Riddell, Rockmore, & Vana, 2023; Reisenbichler et al., 2022), this research is among the first to demonstrate superhuman perceptual evaluations and marketing effectiveness of synthetic marketing imagery across a comprehensive set of marketing applications and generative text-to-image models. Thereby, we shed light on the new paradigm of *generative marketing*—using generative AI to automate or assist marketing activities—which will likely fundamentally change the creation of marketing content in the future.

Second, our findings deepen the understanding of the human perception of AI-generated content. Studies on human perception of advertising content have a long tradition in the marketing literature (e.g., MacKenzie et al., 1986; Pieters & Wedel, 2004). However, due to the recent advent of generative text-to-image models, little is known with respect to consumer perception of synthetic marketing imagery. Are AI-generated images only more cost-efficient to produce, or can they attain human-made images' perceptual evaluations? Our study demonstrates that AI-generated images can exceed human quality and aesthetic levels. An AI model specialized in photorealistic images, namely, Realistic Vision, can even create synthetic marketing imagery that humans perceive as more realistic than real images, which is in line with recent findings on "AI hyperrealism" (Miller et al., 2023). In addition, we explore which visual features can explain differential perception of AI-generated imagery. For example, we observe a negative association between the color saturation and all three perceptual dimensions (quality, realism, and aesthetics).

Third, the present paper adds to the rich body of comparative method studies in marketing (e.g., Hartmann, Heitmann, Siebert, & Schamp, 2023; Krugmann & Hartmann, 2024). Despite the remarkable performance of all AI models, we find that model choice matters. While DALL-E 3 and Midjourney v6 constantly rank among the winning models, SDXL Turbo provides inferior performance compared to other AI models and the human-made benchmark images across almost all applications.

2. Related literature

2.1. The importance of marketing imagery

Images are a core component of contemporary marketing communications (Dzyabura et al., 2023) and "worth a thousand words" (Li & Xie, 2020). Their persuasive power is well documented in the field of marketing, explaining their widespread adoption across diverse marketing contexts, such as online and offline advertising (Pieters & Wedel, 2004; Hartmann et al., 2021), social media (Beichert et al., 2024; Li & Xie, 2020), online shopping (Dzyabura et al., 2023; Zhang et al., 2022b), product design (Burnap, Hauser, & Timoshenko, 2023), and visual consumer reviews (Zhang & Luo, 2023). Furthermore, the abundance of visual data allows brand managers to visually "listen in" and derive actionable insights to position their brands (Liu et al., 2020; Dzyabura & Peres, 2021).

What explains the popularity and power of visual content in marketing communications? Ample evidence supports the picture superiority effect, whereby images are remembered better than words (Childers & Houston, 1984; Paivio & Csapo, 1973). Pieters and Wedel (2004) find that the pictorial element of print ads is superior in capturing attention. Li and Xie (2020) demonstrate a positive mere presence effect of image content on social media engagement compared to text-only content. Their potential to shape consumers' cognitive, emotional, and behavioral responses makes images an appealing medium to convey a brand's visual identity in a memorable way and call consumers to action (Phillips, McQuarrie, & Griffin, 2014).

However, producing visual content is a resource-intensive process. Consider the development of a multi-modal digital ad consisting of a visual component and a corresponding tagline. Each component demands a specialized skill set to create effective content. While the textual tagline can be efficiently revised, the visual element adds a layer of complexity, which arises not only from the aesthetic decisions involved in the creation process (Zhang et al., 2022b) but also from the technical demands of dedicated image-editing software. What if generative AI could support the efficient production of both synthetic textual marketing content (e.g., Reisenbichler et al., 2022; Reisenbichler et al., 2023) and of visual marketing content?

2.2. Generative AI-enabled productivity gains and cost savings

Both within and outside of the marketing context, there is a growing interest in studying generative AI-enabled productivity gains and cost savings. Brynjolfsson et al. (2023) show that access to a generative AI-based conversational assistant can increase customer support agents' productivity by 14%, on average, with even larger benefits for novice and low-skilled workers. For search engine optimization (SEO) content generation, Reisenbichler et al. (2022) demonstrate that large language models can produce significantly more effective text than human SEO experts, quasi-experts, and novices while simul-

taneously incurring a cost benefit of 91%. Similarly, [Reisenbichler et al. \(2023\)](#) show that machine-written ad content can increase the production efficiency of search engine advertising (SEA) by more than 60%.

Beyond marketing, [Dell'Acqua et al. \(2023\)](#) demonstrate for 18 realistic consulting tasks that consultants with GPT-4 access, on average, completed 12.2% more tasks and were 25.1% faster compared to a control group. Similar to the findings of [Brynjolfsson et al. \(2023\)](#), consultants with below-average skills benefited more from generative AI support. [Peng, Kalliamvakou, Cihon, and Demirer \(2023\)](#) present evidence that software developers with access to GitHub Copilot, an AI pair programmer, completed a programming task 55.8% faster than a control group without access. Similarly, [Zhou and Lee \(2024\)](#) show that generative AI can increase human creative productivity by 25% in the context of digital artworks.

Collectively, recent research provides converging evidence for the substantial productivity gains and cost savings enabled by generative AI across various application contexts and data modalities. These findings are in line with the priorities of the business sector, where "tactical benefits such as improving efficiency/productivity (56%) and/or reducing costs (35%)" are reported as primary objectives, according to a survey among 2,835 business and technology leaders ([Ammanath et al., 2024](#)). In addition, there seems to be a growing consensus that generative AI exerts an equalizing effect, narrowing the skill gap between higher- and lower-ability (marketing) content creators (e.g., [Zhou & Lee, 2024](#); [Noy & Zhang, 2023](#); [Zhang, Zhou, & Lee, 2024](#); [Brynjolfsson et al., 2023](#); [Dell'Acqua et al., 2023](#)). However, the adoption of generative AI as a new technology is a function not only of efficiency but also of effectiveness gains. Hence, an important question remains: How do consumers react to AI-generated marketing content?.

2.3. Consumer response to AI-generated content

Pioneering work demonstrates favorable consumer response to generative AI in the context of textual online marketing content generation ([Reisenbichler et al., 2022](#)). Similarly, [Carlson et al. \(2023\)](#) show that AI-generated content can be indistinguishable from human-made content in the automated creation of online reviews. These technological capabilities even enable the generation of personalized persuasion at scale ([Matz et al., 2024](#)) with important implications for society ([Jakesch, Hancock, & Naaman, 2023](#)).

Research on consumer reactions to AI-generated visual content shows mixed results. Challenging the prevailing assumption that artistic, creative tasks are reserved for human intelligence ([Feuerriegel, Hartmann, Janiesch, & Zschech, 2024](#)), [Zhang et al. \(2024\)](#) demonstrate that AI-assisted artists are more likely to receive positive reactions to their creative content. In contrast, [Horton, White, and Iyengar \(2023\)](#) demonstrate that people devalue AI-generated art even if they cannot distinguish it from human-made art. [Jansen et al. \(2024\)](#) show that consumer responses can inform an automated alignment process, whereby a generative AI model is steered to convey certain brand dimensions in its outputs. Also in the context of product design, AI-generated visuals can be appealing to consumers ([Burnap et al., 2023](#); [Zhang, Bai, & Ma, 2022a](#)). However, AI-assisted product designs can also backfire when used in the wrong context ([Xu & Mehta, 2022](#)).

3. Overview of studies

To explore consumer response to AI-generated visual marketing content across various industries and application contexts, we conduct three studies. [Table 1](#) summarizes the objective and setup for each of them, exploring generative AI's potential to rival human-made content both in the lab and in the field ([van Heerde, Moorman, Moreau, & Palmatier, 2021](#)).

First, study 1 compares 10,320 AI-generated marketing images to 2,400 human-made images on three perceptual dimensions: quality, realism, and aesthetics. Study 2 gives identical creative briefings to both the AI models and commissioned human freelancers, allowing us to quantify image production costs and obtain a broader set of standard marketing metrics in an advertising pretest setting. Specifically, we assess image performance along image perception, attitude, behavioral intention, and prompt following, i.e., compliance of the created images with the initial creative briefing. Building on studies 1 and 2, study 3 is designed to investigate the real-world marketing effectiveness of AI-generated banner ads via an actual online marketing campaign featuring seven synthetic images and a human-made stock image selected by an online marketing professional. As the behavioral response, we assess each image's CTR based on 173,022 impressions and 907 clicks. Collectively, the three studies comprehensively assess AI-generated visual marketing content based on over a quarter million human evaluations for 10,355 synthetic marketing images.

All studies draw on the same set of seven state-of-the-art generative text-to-image models that reflect a heterogeneous collection of AI models with differing characteristics such as training data, source accessibility, and release date. [Table 2](#) presents an overview of the AI models included in our studies. [Web Appendix A.1](#) describes the models from a technical perspective. The earliest model we include in our benchmark dates back to October 10, 2023 (Adobe's Firefly 2). The latest model is from February 1, 2024 (Google's Imagen 2).

4. Study 1: Perception of AI-generated vs. human-made marketing imagery

The objective of study 1 is to investigate if AI-generated images can achieve similar perceptual ratings as human-made images in terms of quality, realism, and aesthetics across common visual marketing applications. All three perceptual dimensions are frequently studied image characteristics in marketing research with important downstream consequences ([Zhang](#)

Table 1
Overview of studies.

	Study 1	Study 2	Study 3
Research question	How is AI-generated marketing imagery perceived compared to human-made images across a broad range of real-world marketing applications?	Given the same prompt, i.e., creative briefing, can generative text-to-image models reach similar performance across key marketing outcomes vs. freelancers' human-made imagery?	In a real-world A/B test, can AI-generated images achieve similar click-through rates to a high-quality, human-made stock photo selected by an online marketing professional?
Dependent variables	<ul style="list-style-type: none"> • <i>Quality</i> (Zhang et al., 2022b) • <i>Realism</i> (Cho et al., 2014) • <i>Aesthetics</i> (Talebi & Milanfar, 2018) 	<ul style="list-style-type: none"> • <i>Perception</i> (Zhang et al., 2022b; Cho et al., 2014; Talebi & Milanfar, 2018; Smith et al., 2007) • <i>Attitude</i> (Smith et al., 2007) • <i>Behavioral intention</i> (Smith et al., 2007; Rizzo et al., 2023) • <i>Prompt following</i> (Saharia et al., 2022) 	<ul style="list-style-type: none"> • <i>Click-through rate (CTR)</i>
Study setup	Lab setting (within subjects): <ul style="list-style-type: none"> • 2,400 human-made images • 10,320 synthetic images • 2 × 10 ratings/ image (254,400 ratings in total)	Lab setting (between subjects): <ul style="list-style-type: none"> • 4 human-made images • 28 synthetic images • ~ 50 participants/ condition (1,575 participants in total)	Field study: <ul style="list-style-type: none"> • 1 human-made image • 7 synthetic images • ~ 100 clicks/ condition (173,022 impressions in total)

Table 2
Overview of state-of-the-art generative text-to-image models.

Model	Developer	Source accessibility	Release date	Batch size @ Resolution	Scalable access	Watermark
DALL-E 3	OpenAI	Proprietary	Oct 19, 2023	1@1,024 × 1,024	✓ (API)	~ (invisible C2PA)
Midjourney v6	Midjourney	Proprietary	Dec 21, 2023	4@1,024 × 1,024	✓ (via wrapper API)	X
Firefly 2	Adobe	Proprietary	Oct 10, 2023	4@2,048 × 2,048	X	X (only in free version)
Imagen 2	Google	Proprietary	Feb 01, 2024	4@1,536 × 1,536	X	~ (invisible SynthID)
Imagine	Meta	Proprietary	Dec 06, 2023	1@1,024 × 1,024	X	✓ (see Web Appendix Figure A.2)
SDXL Turbo	Stability AI	Open source (Hugging Face)	Nov 28, 2023	1@512 × 512	✓ (local hosting)	X
Realistic Vision	Community	Open source (Civitai)	Dec 23, 2023	1@512 × 512	✓ (local hosting)	X

Note: DALL-E 3 offers an automated prompt refinement to improve the output. To treat all AI models equally, we followed OpenAI's guidelines to reduce the impact of automated prompt refinement. For details see: <https://platform.openai.com/docs/guides/images/prompting> (accessed March 13, 2024).

et al., 2022b; Li & Xie, 2020; Karpinska-Krakowiak & Eisend, 2024; Kim, Choi, & Wakslak, 2019). Moreover, they represent established evaluation measures for AI-generated images in computer science (Saharia et al., 2022; Betker et al., 2023).

4.1. Method

To ensure comparability between the AI-generated and the human-made images, each synthetic image is created based on a textual description of its underlying human-made source image. Fig. 1 illustrates the two-step image creation process: First, converting human-made source images to text (green trapezoid). Second, generating synthetic sibling images from text (purple trapezoid).⁴

As human-made source images, we utilize 2,400 images that differ by marketing use case and visual composition. To comprehensively reflect prevalent use cases of marketing imagery in our benchmark, we build on established work in marketing, which differentiates between (a) the *data source*, i.e., firms vs. users (Dzyabura et al., 2022; Liu et al., 2020) and (b) the *marketing objective*, i.e., conveying a firm's brand identity as a long-term objective vs. calling prospective customers to action as a short-term objective such as purchasing a product or clicking on an ad (Keller, 1993; Keller & Lehmann, 2006).⁵ This

⁴ Some AI models generate a batch of images for each prompt, e.g., four images for Firefly 2 and Midjourney v6. For these, we consistently sample the first/top left image. To verify the quality of the transformation process, we assess the image-text alignment between the source image and the resulting textual representation as well as between the textual representation and the generated images for a sample of images, following the protocol outlined by Google in Saharia et al. (2022). The resulting image-text alignment is in the range of Google's evaluations, validating our two-step image creation approach.

⁵ The differentiation of the marketing objective in terms of performance marketing (short-term) vs. brand building (long-term) is also a known "marketing dilemma" in the industry (Kyriakidi, 2022).

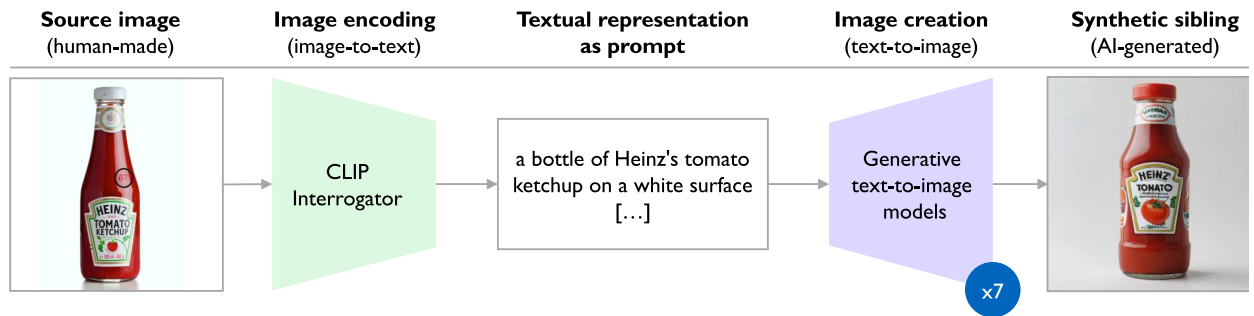


Fig. 1. Image generation procedure in studies 1 and 3. *Note:* In the first step, we employ CLIP-Interrogator in a zero-shot manner via an API endpoint provided by replicate (Ding et al., 2023), allowing us to transform the visual information stored in the human-made source images into textual descriptions (Radford et al., 2021). In the second step, we generate seven synthetic sibling images for each human-made source image using the CLIP-based textual description as a prompt for the seven generative text-to-image models.

conceptual framework results in a 2×2 matrix (see Fig. 2), which structures marketing imagery's most relevant use cases into four quadrants.

For each of the four quadrants, we obtain two representative real-world datasets from which we randomly sample 300 images each.⁶ As acknowledged by (Dzyabura et al., 2022), platforms and specific data sources constantly change, but the general characterization of marketing datasets in terms of the two primary dimensions (firm-generated vs. user-generated and conveying brand identity vs. calling to action) are likely to prevail. Hence, the datasets we sample shall only be considered exemplary for each quadrant. Beyond ensuring diversity in real-world applications and industry contexts, our conceptual framework that guides our systematic data sampling also ensures a large heterogeneity in the human-made images' composition and characteristics. For a detailed overview of the datasets and their provenance, see Web Appendix Table A.1.

To obtain large-scale perceptual ratings for the 2,400 human-made images and their 10,320 synthetic siblings, ten human raters on Amazon Mechanical Turk (MTurk) assessed each image on 7-point Likert scales (1 = lowest, 7 = highest). Before human evaluation, each image was resized to a standardized resolution with a minimum dimension of 512 pixels in either width or height to avoid image distortions or cropping (Sauer, Lorenz, Blattmann, & Rombach, 2023). We adopt the image quality scale from Zhang et al. (2022b): "Give a score to an image on a scale of 1–7 on its aesthetic quality where 1 is "very bad" and 7 is "excellent".". In addition, we provide the same detailed instructions as Zhang et al. (2022b) to ensure reliable ratings (see Web Appendix Figure A.3). To assess perceived realism, we adopt a scale item from Cho, Shen, and Wilson (2014), which Karpinska-Krakowiak and Eisend (2024) also use in the context of deepfake content: "The visual elements of the ad are realistic.", anchored by "strongly disagree" and "strongly agree". Below the question, we define realistic as "accurately representing what is natural or real" (see Web Appendix Figure A.3).⁷

This results in a total of 254,400 human ratings $((2,400 + 10,320) \text{ images} \times 2 \text{ questions/image} \times 10 \text{ raters/question})$. In addition, to enrich these human evaluations with an algorithmic aesthetics assessment, we apply Neural Image Assessment (NIMA) to all our images. NIMA is a convolutional neural network-based classifier for human perception of image aesthetics, defined as $x_{NIMA} \in [1, 10]$, where 10 is the highest score (Talebi & Milanfar, 2018). Web Appendix Figure A.4 plots the highest vs. lowest rated AI-generated images across quality, realism, and aesthetics. The lowest-realism image, for example, features an incoherent object representation, showing a baseball player with an ill-positioned arm.

4.2. Results

Table 3 displays the OLS regression results for the three perceptual dimensions. To account for unobserved differences across the prompts and human raters, we include fixed effects for both. The low to medium correlations among the three dependent variables suggest that they capture three distinct visual characteristics of the images (see Web Appendix Figure A.5).

Regarding image quality, four out of the seven AI models significantly surpass the human-made images, with Firefly 2 ($\beta_{Quality} = .1564, p < .001$) generating the highest-quality images. In contrast, SDXL Turbo obtains the lowest quality perception scores, exhibiting inferior performance compared to the human benchmark ($\beta_{Quality} = -.1268, p < .001$). However, given that ratings are captured on 7-point Likert scales, with an overall mean quality rating of 5.26, the effect sizes suggest only marginally lower quality evaluations than the real-world marketing imagery. DALL-E 3's images, for example, are rated less than four hundredths of a Likert-scale point lower than the human-made benchmark images ($\beta_{Quality} = -.0366, p < .001$).

⁶ For those models that we could not access programmatically, e.g., via an API endpoint, we sampled 30 human-made images per dataset, resulting in a total of 240 images per model for evaluation. This corresponds to 10% of the data compared to the other models we evaluate on 2,400 images each.

⁷ In study 2 we use the full six-item scale for perceived realism by Cho et al. (2014) (Cronbach's $\alpha = .88$).

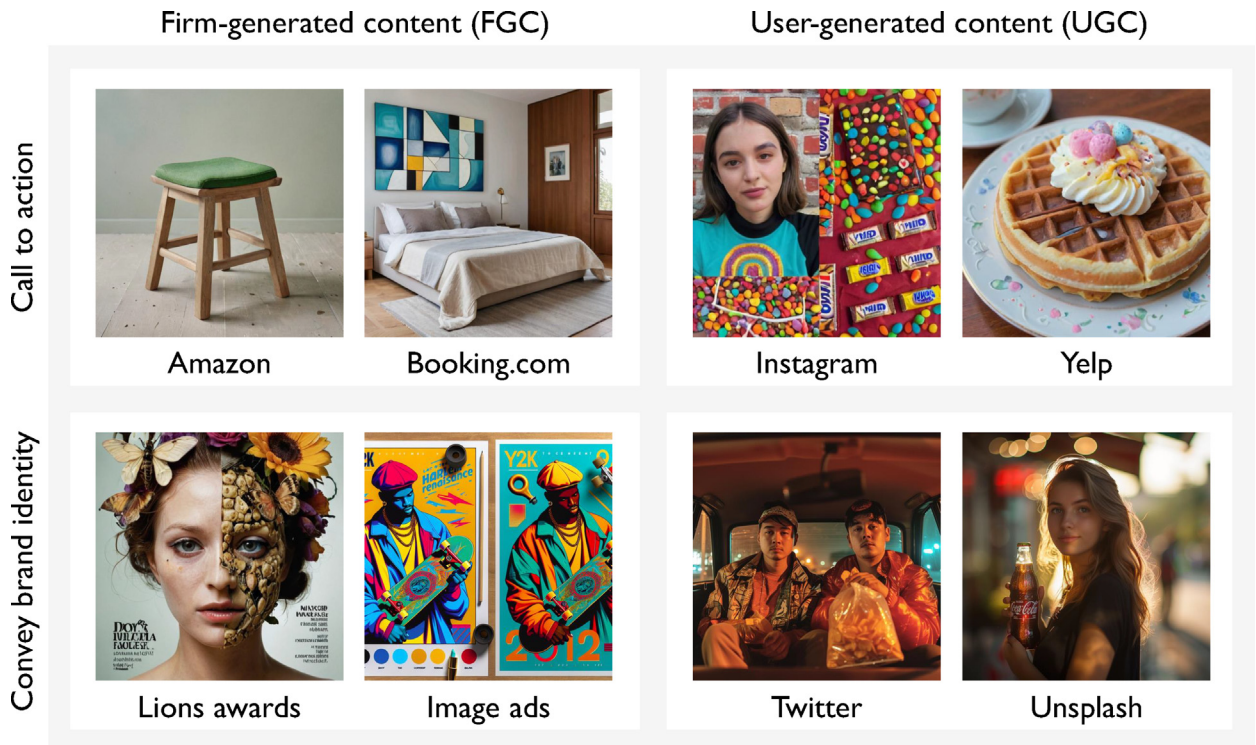


Fig. 2. Overview of datasets. Note: Each quadrant contains two representative datasets. All displayed images are AI-generated.

Realistic Vision lives up to its name, standing out as the only AI model that outperforms the human-made images in terms of perceived realism ($\beta_{Realism} = .0708, p < .01$). This remarkable performance can be attributed to its specialized fine-tuning to generate highly realistic images (see Web Appendix A.1 for details). In contrast, all other AI models fail to achieve the realism level of human-made images.

Regarding aesthetic appeal, all seven AI models significantly surpass the human-made marketing images ($p < .001$). Overall, Firefly 2 and Realistic Vision produce the images with the highest aesthetic appeal.

4.3. Discussion

Study 1 demonstrates that the best AI models can outperform human-made marketing visuals across three important perceptual dimensions: quality, realism, and aesthetics. Furthermore, the results suggest that model choice matters and depends on the advertiser's objective. Regarding image quality and aesthetics, Firefly 2 emerges as the winning method. In terms of perceived realism, Realistic Vision's synthetic images are perceived as more realistic than real, human-made images. This observation is consistent with findings by Jakesch et al. (2023), demonstrating that AI-generated text can be perceived as "more human than human". Similarly, Miller et al. (2023) document so-called "AI hyperrealism" for the perception of AI-generated human faces. For advertisers, these findings are important as realistic product portrayals can facilitate consumers' mental simulation of product consumption or usage (Kim et al., 2019), which in turn can translate into purchase intention (Ceylan, Diehl, & Wood, 2024).

In terms of the mechanism – that is, the images' visual features correlating with differential perceptual assessments – our analysis identifies several visual ingredients that shape consumers' perception of AI-generated marketing imagery (see Web Appendix Table A.3 for details). For example, excessive color saturation appears to diminish favorable consumer responses. Similarly, content creators need to exercise caution when prompting AI models to generate text or human faces, as these elements are prone to visual imperfections, which in turn can hamper the synthetic images' appeal.

Next, study 2 addresses three limitations of study 1. First, study 1 assesses only perceptual dimensions, which we expand in study 2 to a broad set of marketing metrics commonly used advertising pretests, including ad attitudes and purchase intentions (MacKenzie et al., 1986). Second, study 2 simplifies image creation through an alternative prompting procedure. Instead of the two-step pipeline (first, converting images to text; second, generating images from text; Fig. 1), both AI models and commissioned human freelancers receive the same creative briefing, streamlining the process to a text-to-image task. Third, we could not observe the production costs of the human-made images in study 1. Study 2 allows us to obtain these

Table 3
Results for perception of AI-generated vs. human-made marketing imagery.

Dependent variables: Model:	Perception		
	Quality (1)	Realism (2)	Aesthetics (3)
<i>AI models</i>			
DALL-E 3	-.0366*** (.0090)	-.5273*** (.0103)	.1602*** (.0129)
Midjourney v6	.0925*** (.0090)	-.1142*** (.0103)	.0704*** (.0129)
Firefly 2	.1564*** (.0228)	-.0909*** (.0261)	.3994*** (.0326)
Imagen 2	.0156 (.0227)	-.2365*** (.0260)	.2798*** (.0326)
Imagine	.0589** (.0227)	-.2229*** (.0260)	.2880*** (.0326)
SDXL Turbo	-.1268*** (.0090)	-.2945*** (.0103)	.0983*** (.00129)
Realistic Vision	.0960*** (.0090)	.0708*** (.0103)	.2808*** (.0129)
<i>Fixed effects</i>			
Prompt Respondent	Yes	Yes	Yes
	Yes	Yes	No
<i>Fit statistics</i>			
Observations	127,200	127,200	12,720
R ²	.3742	.4003	.5567
Within R ²	.0074	.0348	.0584

Note: Standard errors in parentheses. Human-made image as reference. For aesthetics, NIMA provides a single deterministic prediction whereas ten respondents rated each image's quality and realism. This explains the ten-fold difference in observations for Models (1–2) vs. Model (3).
***: $p < .001$, **: $p < .01$, *: $p < .05$.

costs and compute a back-of-the-envelope calculation on the cost savings for generative AI vs. human labor for visual content creation.

5. Study 2: Generative AI vs. human freelancers

Study 2 is designed to investigate how AI models perform compared to experienced human freelancers. To set the generative AI models and the designers on equal footing, we instruct both with the identical creative briefing and do not share any additional information. Fig. 3 displays a schematic overview of the image creation process. Compared to study 1, this approach is significantly more costly and less scalable. However, the smaller number of images allows us to obtain consumer responses to a broader range of important marketing metrics, thereby mimicking a "diagnostic pretesting" setting (MacKenzie et al., 1986).

5.1. Method

Guided by our framework introduced in study 1 (see Fig. 2), we define one creative briefing for each quadrant of our 2×2 matrix covering prevalent marketing imagery applications (data source: firm-generated content (FGC) vs. user-generated content (UGC); marketing objective: call to action vs. convey brand identity). Specifically, we include a marketing decal showcasing a new limited edition product (i.e., FGC; call to action), a consumer brand selfie (i.e., UGC; call to action), a large outdoor banner ad (i.e., FGC; convey brand identity), and a consumer action shot with merchandise (i.e., UGC; convey brand identity). We choose "Red Bull" as an exemplary brand due to its frequent examination in marketing research (e.g., Seiler, Tuchman, & Yao, 2021), being renowned for its "buzz marketing" (Steenkamp, van Heerde, & Geyskens, 2010). Fig. 4 presents the creative briefings alongside the AI-generated images by DALL-E 3 and human-made images.⁸

To obtain human-made images from experienced human freelancers, we created an individual request for proposal (RFP) for each creative briefing on Freelancer.com and waited until we received more than 50 bids. Based on expert input from an

⁸ The creative briefings are formulated in a similar style to standard prompts, ensuring they are compatible with all seven AI models without modifications and human freelancers can easily understand them.

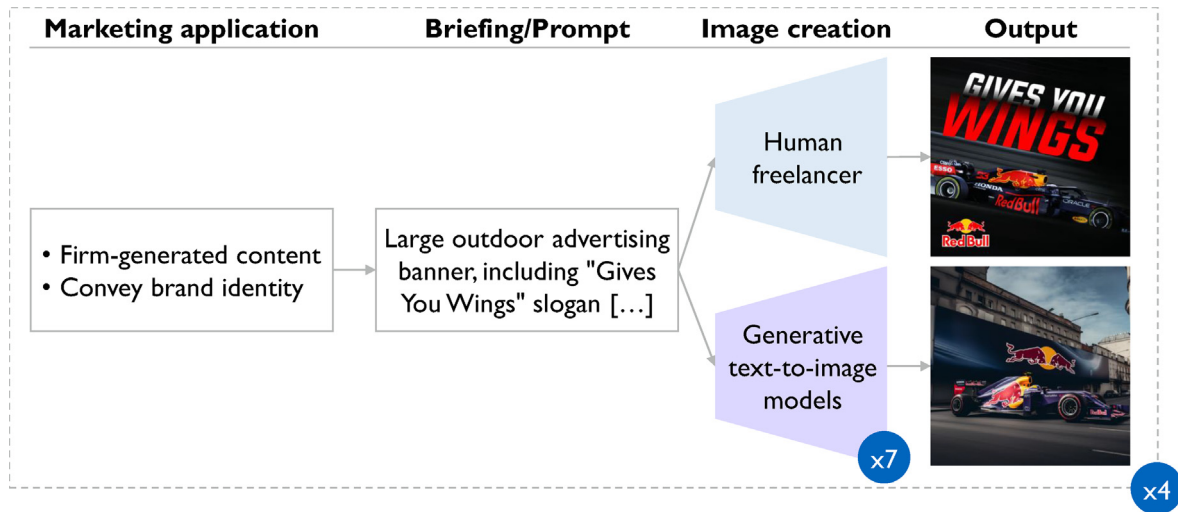


Fig. 3. Image generation procedure in study 2. *Note:* For complete briefing/prompt, see Fig. 4. The blue circles indicate that we use seven generative text-to-image models to obtain AI-generated sibling images for four different marketing applications.

independent freelancer unaware of our hypothesis, we defined a budget range of USD 30 to USD 250 per image. Besides the creative briefing, the RFP included additional information, such as two statements that prohibit the use of any generative AI tools⁹ as well as the desired quadratic aspect ratio of the output (see Web Appendix Figure A.6 for one of the four RFPs used to brief the freelancers). Among the bids, we then filtered for freelancers that (a) commit in writing not to use any generative AI tools, (b) have a rating of 4.5 out of 5 stars or higher, (c) have a bid above the average bid, and (d) are verified by Freelancer.com. This systematic selection results in three unique freelancers offering their service at USD 100 per image. Their profiles state hourly rates ranging from USD 15 to USD 30.

Based on the same creative briefings shared with the human freelancers, we prompt all generative text-to-image models. This approach results in a total of 32 images where four are human-made and 28 are AI-generated (i.e., 7 AI models \times 4 images). We use these images as stimuli for a between-subjects experiment on Prolific, where each participant is exposed to only one of the 32 images. We collect ten dependent variables, which can be subsumed into the following four groups:

- Perception via *quality, realism, aesthetics, and ad creativity*
- Attitude via *ad attitude and brand attitude*
- Behavioral intention via *purchase intention and social media engagement*
- Prompt following via *image-text alignment and brand recognition*

All questions are rated on 7-point Likert scales (1 = lowest, 7 = highest), using established multi-item scales. Web Appendix Table A.5 gives a detailed overview of these scales and respective references from the marketing literature.

5.2. Results

1,575 of the 1,604 Prolific panelists pass all three attention checks, resulting in an average of 49 participants per condition ($M_{age} = 43.35$ years, 50.16% women). Table 4 presents the OLS regression results.

Overall, DALL-E 3 produces the best synthetic images, significantly outperforming the human freelancers in terms of five marketing metrics, and obtaining directionally higher evaluations across the other five. Midjourney v6 ranks second, surpassing the freelancers on four evaluation criteria, and being directionally better on the other ones, except for brand recognition. The open-source models, SDXL Turbo and Realistic Vision, exhibit the worst performance and are almost consistently inferior to the freelancers' marketing imagery.

Regarding perceptual dimensions, participants rate the aesthetic appeal of all AI-generated images significantly higher than the freelancers' visuals, which aligns with our observation in study 1. Also for quality and realism, the best AI models outperform the human-made images (DALL-E 3 and Midjourney v6 for quality, and Midjourney v6 and Imagen 2 for realism). DALL-E 3 is the only model that obtains significantly higher assessments of ad creativity ($\beta_{AdCreativity} = .4198, p < .01$), while

⁹ This restriction on AI use mirrors current trends in the formulation of agency contracts (Sloane, 2024).

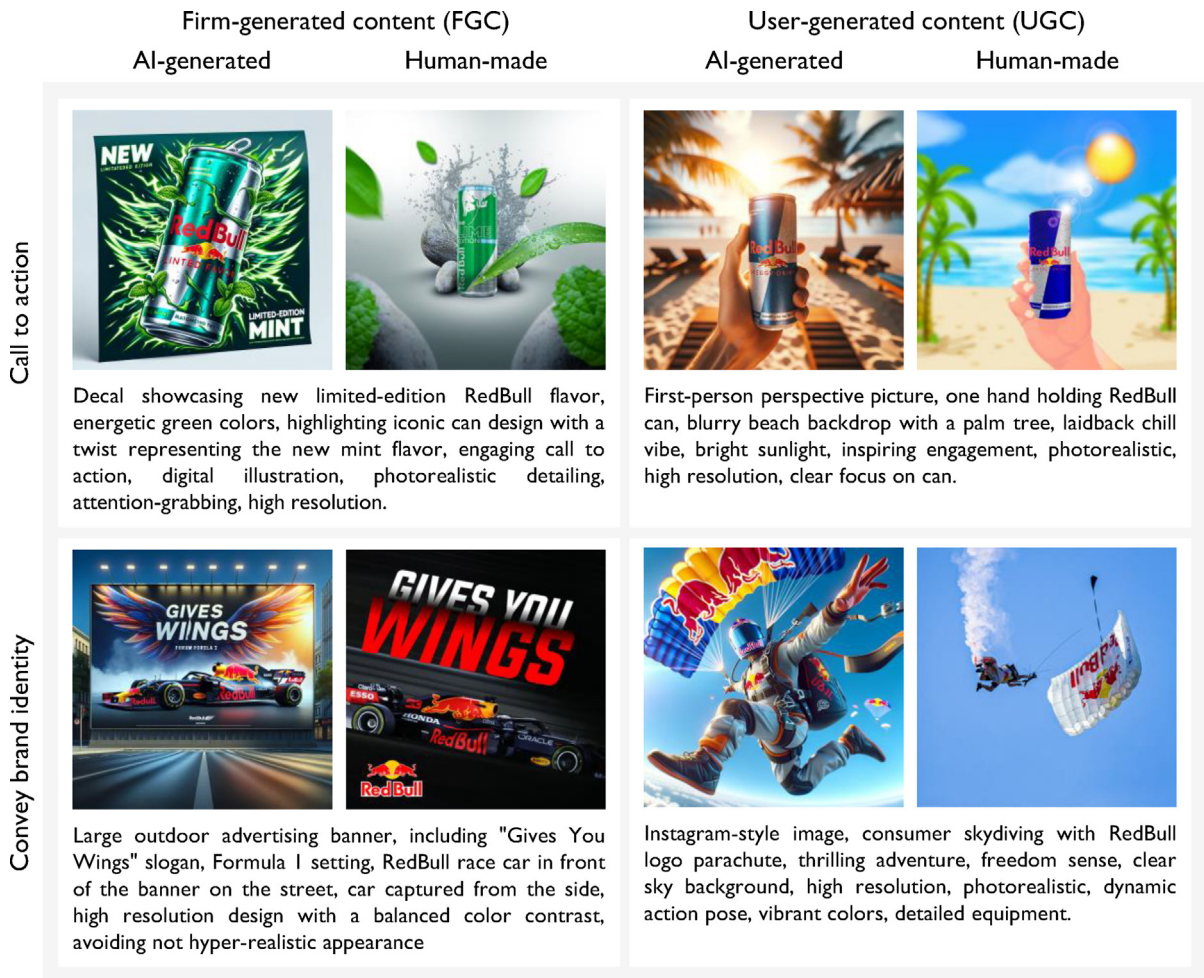


Fig. 4. Overview of four creative briefings and resulting images. Note: DALL-E 3 created the AI-generated images (left). Freelancers created the human-made images (right).

SDXL Turbo and Realistic Vision exhibit significantly lower ad creativity compared to the freelancers ($\beta_{AdCreativity} = -.5582, p < .001$ and $\beta_{AdCreativity} = -.7357, p < .001$, respectively).

In terms of attitudinal response and behavioral intentions, DALL-E 3 and Midjourney v6 are the only two AI models that obtain directionally higher assessments compared to the freelancers, outperforming the human-made images significantly regarding ad attitudes ($\beta_{AdAttitude} = .3763, p < .05$, and $\beta_{AdAttitude} = .3901, p < .05$, respectively).

Lastly, DALL-E 3 obtains the highest text-to-image alignment ($\beta_{Alignment} = .5607, p < .001$), capturing participants' response to the question: "How accurately does the caption describe the above image?". In other words, the best AI model manages to generate images that adhere better to the creative briefing than the commissioned human freelancers. Imagen 2 ranks second, also obtaining a significantly higher image-text alignment than the human-made images ($\beta_{Alignment} = .3862, p < .001$). These findings are plausible as both DALL-E 3 and Imagen 2 are trained on enhanced, synthetic image-text pairings to improve their prompt following capability (see Web Appendix A.1).

The open-source models, SDXL Turbo and Realistic Vision, as well as the proprietary Firefly 2 are inferior to the human-made images, both in terms of image-text alignment and brand recognition. This is plausible as Firefly 2 does not generate brand logos at all and the open-source models generate Red Bull logos, which tend to be distorted or highly corrupted, making participants less likely to correctly recognize them as Red Bull (see Web Appendix Figure A.7 for example images).

Table 4

Results for perception, attitude, behavioral intention, and prompt following of AI-generated vs. human-made marketing imagery.

Dependent variables: Model:	Perception				Attitude		Behavioral intention		Prompt following	
	Quality (1)	Realism (2)	Aesthetics (3)	Ad creativity (4)	Ad attitude (5)	Brand attitude (6)	Purchase intention (7)	Engagement (8)	Alignment (9)	Brand recognition (10)
<i>AI model</i>										
DALL-E 3	.6462*** (.1422)	.0854 (.1071)	.6273*** (.0330)	.4198** (.1446)	.3763* (.1641)	.0939 (.1566)	.0078 (.1879)	.1313 (.1870)	.5607*** (.1547)	.0050 (.0280)
Midjourney v6	.6335*** (.1421)	.3843*** (.1070)	.4891*** (.0330)	.1667 (.1445)	.3901* (.1640)	.2070 (.1565)	.2153 (.1878)	.1186 (.1869)	.1540 (.1546)	-.0519 (.0280)
Firefly 2	.1829 (.1449)	-.3178** (.1091)	.5515*** (.0336)	.0560 (.1474)	.0447 (.1672)	-.0279 (.1596)	.1931 (.1915)	.2912 (.1906)	-1.371*** (.1577)	-.9548*** (.0285)
Imagen 2	-.2347 (.1431)	.3965*** (.1077)	.6277*** (.0332)	-.2490 (.1455)	.0963 (.1651)	-.0954 (.1576)	-.2010 (.1891)	-.1854 (.1882)	.3862* (.1557)	-.0479 (.0282)
Imagine	.0897 (.1419)	-.0002 (.1068)	.3433*** (.0329)	-.0855 (.1443)	-.0758 (.1637)	-.0503 (.1562)	-.2597 (.1875)	-.1829 (.1866)	.2047 (.1544)	-.0254 (.0279)
SDXL Turbo	-.8379*** (.1428)	-.6207*** (.1076)	.4900*** (.0331)	-.5582*** (.1453)	-.8801*** (.1649)	-.7553*** (.1573)	-.5985** (.1888)	-.2199 (.1879)	-1.396*** (.1555)	-6.751*** (.0281)
Realistic Vision	-.7618*** (.1431)	-.3223** (.1077)	.6665*** (.0332)	-.7357*** (.1455)	-.8831*** (.1651)	-.5189* (.1576)	-.2696 (.1891)	-.1630 (.1882)	-.8023*** (.1557)	-.3196*** (.0282)
<i>Fixed effects</i>										
Prompt	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>										
Observations	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575
R ²	.1436	.1278	.5559	.1136	.0992	.1002	.1622	.0730	.2396	.6227
Within R ²	.1314	.0939	.2797	.0720	.0938	.0987	.1565	.0707	.1933	.6066

Note: Standard errors in parentheses. Human-made image as reference. All regression models control for the participants' brand familiarity with Red Bull. ***: $p < .001$, **: $p < .01$, *: $p < .05$.

5.3. Quantifying generative AI-enabled cost savings compared to human freelancers

Observing the fees paid to the commissioned human freelancers allows us to compare these with the production costs incurred by the AI models. Specifically, we compute a back-of-the-envelope calculation to determine the cost per image for each of the seven AI models and the freelancers. For details on the underlying input parameters and assumptions, refer to Web Appendix Table A.6.

For those AI models offered via an API endpoint with a transparent pricing scheme the cost per image is straightforward to obtain, e.g., for DALL-E 3 at USD .04 per image or Imagen 2 at USD .02 per image. Other AI models are priced based on the usage of Graphics Processing Unit (GPU) hours or packaged image credits, which translate into costs of USD .07 for Midjourney v6 and USD .05 for Firefly 2. For the open-source models, SDXL Turbo and Realistic Vision, which can be self-hosted on local or cloud GPUs, we obtain the lowest costs at USD .00005 and USD .00026 per image, respectively. In other words, creating a single image with SDXL Turbo costs only one two-hundredth of a cent ($\frac{5}{1000} = \frac{1}{200}$). This is reasonable as SDXL Turbo is optimized for rapid inference through adversarial diffusion distillation, enabling image creation in just a single diffusion step (Sauer et al., 2023). The ability to create nearly 1,000 synthetic images with SDXL Turbo for the same budget compared to generative text-to-image models with an API endpoint such as DALL-E 3, highlights the substantial cost variations among AI models. Note in addition to their cost advantage and inference speed (Sauer et al., 2023), open-source software and models are appealing in terms of control and customizability (Bonaccorsi & Rossi, 2003).

The cost-efficiency advantage of AI models over human freelancers is even stronger. For the human freelancers' images, we paid an average USD 100 per image. In contrast, the same budget allows for the creation of 2,500 images with DALL-E 3 or two million images with SDXL Turbo, showcasing the disruptive efficiency gains offered by generative AI for the automated creation of visual marketing content.

5.4. Discussion

Study 2 investigated the effectiveness and efficiency of AI-generated marketing imagery compared to commissioned human freelancers' work in a between-subjects setting, mimicking an advertising pretest. DALL-E 3 emerges as the winning generative text-to-image model, significantly outperforming the human-made visuals in 50% of the marketing metrics, with directionally higher ratings on the other 50%. Notably, DALL-E 3 is the only AI model to achieve significantly higher ad creativity ratings compared to the freelancers' visuals. Midjourney v6 ranks as the second most effective AI model. The findings align with the disruptive impact of generative AI on creative freelancer work, as suggested by Hui, Reshef, and Zhou (2023). While not without flaws, generative AI's output is often perceived as comparable to, if not better than, human-made content,

but at a substantially lower cost. Strikingly, the AI models generated images that closely matched Red Bull's iconic design elements, such as the distinctive shape of the can and the brand's characteristic color palette (see Fig. 4). This suggests that even without brand-specific fine-tuning the AI models had sufficient examples within their extensive training datasets to learn and replicate Red Bull's visual language.

A common issue for both developers and users of generative AI is ensuring that AI models follow instructions precisely, producing outputs that match the given prompts (Betker et al., 2023). Study 2 suggests that this concern may extend to commissioned human freelancers who sometimes deviate from the creative briefing provided to them. For example, despite the request for a *mint*-flavored limited edition Red Bull can, one freelancer created a *lime*-flavored design, see Fig. 4. Similarly, another freelancer was instructed to create a "photorealistic" and "high resolution" brand selfie but delivered an illustration instead. These instances highlight that even careful freelancer selection does not guarantee flawless results and that some images might be more complex to create for humans than others. While iterations with the freelancers could have corrected some of these flaws, we did not do so in order to maintain a consistent one-shot creation process for both the AI-generated and the human-made images. Note that had we opted for iterations, we could have iterated the marketing visuals much faster with the AI models than with the human freelancers.

Next, study 3 focuses on the real-world performance of AI-generated banner ads in a field study designed to collect additional evidence on the effectiveness and efficiency gains enabled by generative AI.

6. Study 3: The real-world effectiveness of AI-generated banner ads

To increase the ecological validity of our findings (van Heerde et al., 2021; Hulland & Houston, 2021), study 3 systematically assesses the real-world effectiveness of synthetic images compared to a professional human-made stock photo. Specifically, we measure performance in terms of the AI-generated banner ads' CTR, i.e., clicks divided by impressions.

6.1. Method

For the field study, we collaborated with an education provider specializing in online marketing courses. To mimic a real-world marketing campaign as closely as possible, we co-designed the experiment with the education provider's CEO, an online marketing expert holding a Ph.D. in marketing with 25 years of industry experience.

First, the online marketing expert selected and purchased the human-made stock photo in line with the education platform's marketing objective. The image displays two hands holding two puzzle pieces against a horizon at sunset (see Fig. 5). Note that hands are considered notoriously difficult to create for generative AI models (Chayka, 2023). Hence, the online marketing expert did not choose an easy baseline for the generative text-to-image models to compete with. Moreover, the selected stock photo originates from an iStock account that offers more than 4,500 images in its portfolio (Oatawa, 2024), suggesting professional experience in the creation of stock photography.

To generate seven synthetic sibling images, one per AI model, we followed our validated image creation approach from study 1 (see Fig. 1). All banner ads had the same title ("Understand online marketing with [company name].") and a detailed description ("Learn to piece together the components for your online marketing success! With [company name] - your expert for online marketing training. We train on over 70 online marketing topics in small groups!").

To ensure fair competition between the banner ads, we chose Meta's online marketing platform which offers a sophisticated A/B testing functionality.¹⁰ As Meta restricts the number of ads that can be included in a single campaign when the A/B testing functionality is activated, we created two concurrent A/B testing campaigns that both include the professional stock photo as the human-made benchmark. Both campaigns had identical campaign settings, targeting an audience in the 25–55 age group interested in marketing. In addition, we set "traffic" as the campaign objective to avoid the unobservable Meta algorithm optimizing the ads' distribution on our dependent variable, namely, clicks. We set a daily budget of USD 10 per condition and aimed for an expected 100 clicks per condition, i.e., 900 clicks in total. Both campaigns started on February 27, 2024, and ran until March 2, 2024.

6.2. Results

The campaigns cost a total of USD 449.91, generating 173,022 impressions and 907 clicks. This translates into an average CTR of .52%. Table 5 reports the CTRs for all banner ads.

DALL-E 3 generated the best-performing banner ad, obtaining a CTR of .80% and significantly outperforming the professional stock photo within the same A/B testing campaign by more than 50% ($CTR_{StockPhoto} = .53\%$; $\chi^2(1, N = 35,110) = 9.5915, p < .01$). Given the same budget for both conditions of USD 49.99 this results in an over 34% higher cost per click ($CPC_{DALL-E3} = USD.38$ vs. $CPC_{StockPhoto} = USD.51$). In addition, we find that model choice matters. Specifically, DALL-E 3 significantly outperforms the worst-performing model SDXL Turbo ($\chi^2(1, N = 36,594) = 23.968, p < .001$) by 100% ($CTR_{DALL-E3} = .80\%$ vs. $CTR_{SDXLTurbo} = .40\%$).

¹⁰ See <https://www.facebook.com/business/ads/ab-testing> for details (accessed March 14, 2024).



Fig. 5. AI-generated banner ads and professional stock photo for field study. *Note:* CLIP-Interrogator transformed the professional stock photo into the following prompt that served as input to all seven generative text-to-image models: “a person holding two pieces of a puzzle, a stock photo by [artist], shutterstock contest winner, objective abstraction, stockphoto, stock photo, congruent”.

Table 5
CTRs of professional, human-made stock photo vs. AI-generated banner ads.

Rank	Model	CTR	Impressions	Clicks	CPC	Spend	Campaign
1	DALL-E 3	.80%	16,579	133	\$.38	\$ 49.99	A
2	Midjourney v6	.54%	19,310	105	\$.48	\$ 49.99	A
3	Imagine	.54%	19,851	107	\$.47	\$ 49.99	A
4	Stock Photo	.53%	18,531	98	\$.51	\$ 49.99	A
5	Stock Photo	.52%	19,606	101	\$.49	\$ 49.99	B
6	Imagen 2	.51%	19,170	97	\$.52	\$ 49.99	B
7	Firefly 2	.49%	19,612	97	\$.52	\$ 49.99	B
8	Realistic Vision	.43%	20,348	88	\$.57	\$ 49.99	B
9	SDXL Turbo	.40%	20,015	81	\$.62	\$ 49.99	A
Total			173,022	907		\$449.91	

6.3. Quantifying generative AI-enabled cost savings compared to professional stock photos

Again, we quantify the cost savings enabled by generative AI. While the professional stock photo is more than ten times cheaper than a freelancer image (USD 9 vs. USD 100), it is orders of magnitude more expensive than the most effective image in the online campaign generated by DALL-E 3 (USD 9 vs. USD .04). Put differently, an advertiser can create 225 images with DALL-E 3 for the price of a single stock photo.

In addition to the production costs, the CTR of DALL-E 3, which is over 50% higher than that of the stock photo, substantially lowers its CPC by over 25%. Spending the same budget of USD 49.99, DALL-E 3 obtains 133 clicks while the stock photo within the same campaign obtains only 98 clicks. The least effective AI model, SDXL Turbo, achieves only 81 clicks, resulting in a CPC over 63% higher than DALL-E 3.

6.4. Discussion

Study 3 showed in a real-world environment that the best AI models can generate synthetic banner ads that surpass the CTR of high-quality, human-made stock photography selected by an online marketing professional. The CTR increase by more than 50% between DALL-E 3 and the best human-made image is noteworthy, especially as we did not conduct any prompt engineering or fine-tuning of the AI models.¹¹

Closely inspecting the image by DALL-E 3 and the stock photo (see Fig. 5) reveals only a subtle difference in their image composition. While the presence of another person in the stock photo might deter observers from engaging in self-referencing (Hartmann et al., 2021), the first-person perspective of the image generated by DALL-E 3 could encourage observers to imagine themselves holding the puzzle pieces, fostering mental simulation and translating into positive downstream consequences (Ceylan et al., 2024). Furthermore, the physical arrangement of objects is more symmetric in the DALL-E 3 image, which can evoke favorable consumer response (Zhang et al., 2022b).

Advertisers commonly run campaigns with multiple assets on online marketing platforms that redistribute the campaign budget based on each asset's effectiveness (Schwartz, Bradlow, & Fader, 2017). Our findings suggest that generative AI fits well into this A/B testing paradigm, as it allows for creating many visual assets at a fraction of the cost of human-made content with the potential to provide at least the same marketing effectiveness. The adoption of a 'human-in-the-loop' system can improve the effectiveness of AI applications in the field even more (Reisenbichler et al., 2022). This approach involves human experts evaluating the quality of multiple candidate ads before launching a marketing campaign. Alternatively, a specialized predictive AI system, informed by historical performance data, can assume the function of the human experts in an 'AI-in-the-loop' arrangement, resulting in an iterative interplay between a generative AI and a predictive AI model.

7. General Discussion

7.1. Summary

Generative AI represents a new paradigm, fundamentally disrupting the marketing industry (Peres et al., 2023). The present paper demonstrated the effectiveness and efficiency gains that state-of-the-art generative text-to-image models can enable across a broad set of marketing use cases. The AI models' ability to rival human-made content across key marketing metrics suggests that firms may soon find it necessary to embrace generative AI for visual marketing content generation in their day-to-day operations to stay competitive.

Study 1 systematically evaluated consumer perceptions of AI-generated vs. human-made marketing imagery, drawing on 254,400 human evaluations and algorithmic aesthetics assessments. The results showed that the best AI models can generate synthetic sibling images that can significantly outperform their human-made benchmark images in terms of quality, realism, and aesthetics. Strikingly, consumers perceived synthetic images produced by Realistic Vision as more realistic than real images, which has important implications beyond the marketing discipline (Miller et al., 2023; Nightingale & Farid, 2022). Study 2 benchmarked the same AI models with experienced human freelancers, giving both the same creative briefings. The results showed that the best AI models (DALL-E 3 and Midjourney v6) can outperform human-made marketing visuals across a broad battery of marketing metrics at a fraction of the cost. Study 3, a field study with over 170,000 impressions, provided evidence on the real-world effectiveness of AI-generated banner ads. DALL-E 3, the best AI model, yielded an over 50% higher CTR than a high-quality, human-made stock photo selected by an online marketing professional. Also, AI model choice matters. Compared to DALL-E 3, the CPC of SDXL Turbo was more than 34% higher. Hence, choosing the wrong AI model can translate into substantial economic costs (Hartmann, Huppertz, Schamp, & Heitmann, 2019).

7.2. Contribution and Implications

This research's large-scale evaluation of AI-generated marketing imagery provides three important contributions for scholars, managers, and policymakers. First, we provide evidence that generative AI can match and even surpass human-made images' consumer perception and marketing effectiveness. While comparative studies are well-established in marketing research (e.g., Andrews, Ainslie, & Currim, 2002; Hartmann et al., 2023), our work, to our best knowledge, is the first systematic comparison between human-made marketing content, such as professional stock photos and visual assets from commissioned human freelancers, and AI-generated images produced by multiple state-of-the-art generative text-to-image models. Our findings aim to assist marketing researchers and practitioners select appropriate AI models for their substantive applications.

Second, the present paper contributes to understanding the human perception of AI-generated visual marketing content. Perceptual studies are important for identifying improvement levers of marketing materials (e.g., Pieters & Wedel, 2004). For

¹¹ These results align with those from a field study we conducted on Taboola's online marketing platform in December 2022 involving 13 generative text-to-image models. In this field study, the top-performing AI model exceeded the human-made stock photo's CTR by over 20%. For illustrative purposes, Web Appendix Figure A.8 presents the human-made benchmark image alongside the AI-generated images that performed best and worst (SDv1-3 and Disco Diffusion, respectively). Note the substantial advancements in image quality observed when comparing the latest generations of AI models to their earlier versions from 1.5 years ago, highlighting the rapid pace of technological progress in automated image creation.

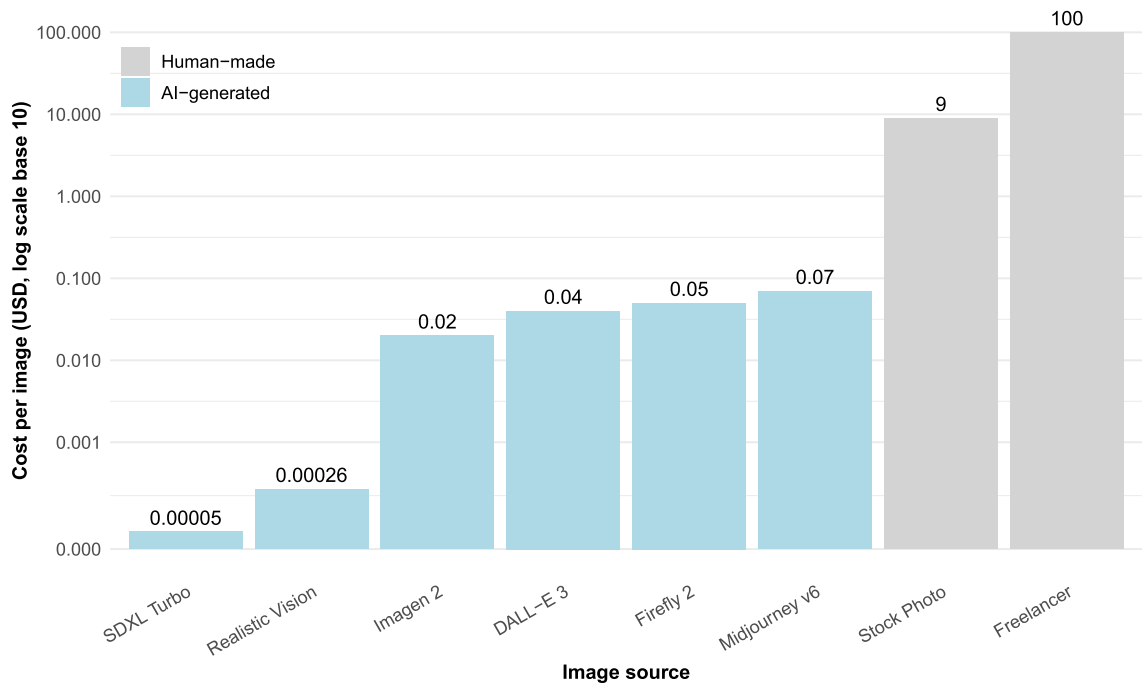


Fig. 6. Cost per image for AI-generated vs. human-made images. *Note:* Cost per image as of their creation dates in Q1 2024.

example, we find that excessive saturation in AI images is negatively associated with perceived quality, realism, and aesthetics. Flaws in human representation can also lead to adverse perceptual reactions by consumers. To facilitate future research on the mechanism between visual features and consumer response, we contribute all our AI-generated images as "GenImageNet" to the research community.¹²

Third, inspired by [Reisenbichler et al. \(2022\)](#), we quantify the productivity gains of generative AI for the creation of visual marketing content. While each freelancer image cost USD 100 (study 2) and the professional stock photo USD 9 (study 3), producing AI-generated images is orders of magnitude more cost-efficient. For example, generating a single image in standard resolution with DALL-E 3, the winning method of our comparative study, costs only USD .04. For the open-source models, SDXL Turbo and Realistic Vision, that can be hosted and run locally, the costs per image converge to zero (USD .00005 and .00029, respectively). [Fig. 6](#) visually summarizes the results of our back-of-the-envelope cost calculations (for details, see Web Appendix Table A.6). Note that the cost per image on the y-axis is log-scaled, highlighting the substantial cost differences across the image sources.

Beyond these three core contributions, our findings are a harbinger of AI-generated images' potential to disrupt visual marketing content generation in the near future. Fitting well into existing marketing paradigms such as rapid A/B testing, where multiple visual assets are evaluated competitively and budget is allocated to the highest-performing ad ([Schwartz et al., 2017](#)), advertisers can expect synthetic images to play a role of growing importance in real-world marketing campaigns. Consistent with our findings, leading online marketing platforms such as Google and Taboola recently announced a seamless integration of generative AI into their ad managers ([Dischler, 2023](#); [Feeney, 2024](#)), enhancing generative AI's accessibility, adoption, and appeal. Fueled by its cost efficiency, the widespread application of generative AI can contribute to more targeted ads with higher individual quality, as slight variations in an advertising message's verbal and visual language can increase its appeal for different target audiences ([Matz, Kosinski, Nave, & Stillwell, 2017](#)). At the same time, this outlook of "personalized mass persuasion" warrants scrutiny by policymakers and scholars across disciplines ([Matz et al., 2024](#)).

What are the implications for companies of different sizes? [Acar and Gvirts \(2024\)](#) suggest that generative AI exerts a leveling effect with "the potential to close the content, insight, and technology gaps that large corporations typically have over their smaller counterparts". The cost and performance advantages that we document across our studies, coupled with the high accessibility of generative text-to-image models, support this notion that generative AI can lead to a democratization of effective visual marketing content. Even if the prediction that "gen AI makes everyone an ad agency" ([Thomas, 2024](#)) does not fully materialize, it will likely substantially reweigh the tasks in the content creation process ([Carlson et al., 2023](#); [Noy & Zhang, 2023](#)).

¹² GenImageNet is available for download at: <https://osf.io/8ctjy/>

Table 6
Examples for research questions related to AI-generated marketing imagery.

Research topic	Exemplary research questions
Advertising	<ul style="list-style-type: none"> • How will AI-generated ads alter the overall creativity of advertising over time? • Can generative AI help in increasing (authentic) diversity in advertising? • How can fine-tuning of generative AI models help in achieving different marketing objectives (e.g., clicks vs. conversions of banner ads)? • How can different data modalities (e.g., image, text, video, audio) be integrated in generative AI-enabled advertising pipelines? • How can the integration of generative AI and A/B testing functionality improve the "learn-and-earn" trade-off? (Schwartz et al., 2017)
Product design	<ul style="list-style-type: none"> • How do usage patterns of generative AI evolve over time for different users groups? • How can generative AI support product design processes, ranging from ideation to hyper-customization for different customer segments? • Which product, brand, or customer characteristics moderate the perception of AI-generated product designs? • Can AI models learn a firm's "visual brand essence" to create on-brand designs?
Social media	<ul style="list-style-type: none"> • Do social media users prefer human or virtual influencers? • What factors moderate consumers' response to virtual influencers? • What is the role of trust in virtual influencers' marketing effectiveness (over time)? • How does generative AI affect online communities and two-sided markets, e.g., freelancer or user-generated artwork? • How do virtual influencers affect consumers' social well-being?
Online shopping	<ul style="list-style-type: none"> • Can generative AI enable personalized visual "website morphing"? (Hauser et al., 2009) • Can AI-enabled product presentations or virtual try-ons reduce product returns? • How might generative AI inflate customer expectations? • How can generative AI enhance the interactivity of online shopping, e.g., via visual assistance or dynamic product presentations such as instant color changes?
Productivity	<ul style="list-style-type: none"> • How will AI-enabled ad makers offered by online marketing platforms disrupt the value chain of content creation, e.g., by commoditizing certain tasks? • How much human supervision is required when using generative AI in fast-paced, multi-asset marketing campaigns? • Can multi-modal models help in emulating customer behavior, e.g., their brand perception by simulating brand elicitation exercises? • What are the benefits and costs of model fine-tuning vs. using commercial off-the-shelf foundation models? • Can zero-shot image analytics using multi-modal large language models replace conventional supervised image classification models?
Moderators	<ul style="list-style-type: none"> • How might new policies and regulation affect the adoption of generative AI in marketing, e.g., enforced disclosure and watermarking of AI-generated content? • Are certain consumer segments (e.g., older, less educated) more vulnerable to AI-enabled deepfakes, misinformation, and deception? • How do consumers react to AI-generated content over time? • What are the implications of biases in generative text-to-image models used for the increasingly automated generation of marketing imagery? • How can potential biases be uncovered and remedied? • How to train AI models while protecting sensitive and private user data? • How should policymakers react to potential threats of personalized mass persuasion?

Lastly, from a societal perspective, our findings have implications for policymakers and contribute to the broader societal debate on the dissemination and detection of deepfakes and disinformation (Karpinska-Krakowiak & Eisend, 2024). Should firms be allowed to use synthetic images in marketing without an AI label? How do consumers react to such a disclosure?¹³ Our results indicate that already today, consumers can perceive synthetic images generated by specialized generative text-to-image models as more realistic than real images. Considering the societal implications of generative AI, early adopters of this disruptive technology must carefully monitor and navigate the rapidly evolving regulatory and legal landscape. Future legislation might, for example, enforce digital provenance standards, such as watermarks and disclosures. Similarly, users of generative AI risk infringing intellectual property rights, requiring firms to exercise caution when selecting an AI model or provider (Inman, Meyer, Schweidel, & Srinivasan, 2024; Feuerriegel et al., 2024; Wang, Bell, Dotson, & Schweidel, 2023).

¹³ Images generated with Meta's AI model, Imagine, include a visible watermark disclosing that they were generated with AI, reading "Imagined with AI". To ensure a fair, undisclosed comparison between all AI models and the human-made benchmark images, we obscured the watermark through a blurring function (see Web Appendix Figure A.2). However, to explore if Meta's AI disclosure relates to human perception of realism, we kept 20% of the images with a watermark (control) and only blurred the watermark on the remaining 80% (treatment) of the images. Based on 2,400 human evaluations of these Imagine-generated images, we find only marginally significant evidence that the disclosure treatment reduces human realism perceptions ($\beta_{\text{watermark}} = -.1241, p = .077$). Apparently, the AI watermark's presence is not sufficient to alter consumer perceptions (see Karpinska-Krakowiak & Eisend (2024) for similar results).

7.3. Limitations and future research directions

We acknowledge that our studies are subject to certain limitations that can inspire future research. While we cover a diverse set of marketing applications, future research can explore the effectiveness of generative AI for additional upper- and lower-funnel outcome measures, e.g., brand awareness or sales. Building on our advertising pretest with commissioned human freelancers from Freelancer.com, future research can benchmark state-of-the-art generative text-to-image models against more expensive professional ad agencies, with and without generative AI access (Sloane, 2024; Thomas, 2024).

Furthermore, by design of our multi-step image creation pipeline (see Fig. 1) and to ensure a fair comparison, we did not conduct any fine-tuning of the AI models. Similarly, we did not conduct prompt engineering to avoid injecting a human bias into the prompt creation. As both of these levers can further enhance generative text-to-image models' effectiveness (Jansen et al., 2024; Rombach, Blattmann, Lorenz, Esser, & Ommer, 2022), our results likely represent a lower bound for the performance of AI-generated marketing imagery. The emergence of future generative AI models will likely improve the synthetic images' perceptual ratings and real-world effectiveness, especially when combined with task-specific data for model calibration (Feng, Zhang, & Srinivasan, 2023).

Lastly, while the presented studies show that generative AI can achieve better results than human-generated content, more research is needed that analyzes the ingredients, i.e., the visual characteristics, that explain consumers' response to AI-generated marketing images (e.g., Zhang & Luo, 2023; Zhang et al., 2022b). While we identified relevant structural and content variables associated with differential perceptual evaluations regarding quality, realism, and aesthetics, future research with larger sample sizes can explore the relationship between additional visual features and lower-funnel effectiveness measures such as conversions and sales. Table 6 lists further research ideas, including risks and moderators that might hamper the adoption of generative AI.

Generative AI fundamentally disrupts visual marketing content generation. This research investigated the disruptive potential of generative AI in marketing in terms of its effectiveness and efficiency. By systematically benchmarking seven state-of-the-art generative text-to-image models against human-made content, we showed that AI-generated marketing imagery can achieve superhuman perceptual evaluations and effectiveness levels in real-world applications at a fraction of the cost of human-made content. We hope our paper inspires future research in the rapidly evolving area of generative marketing.

Data availability

Data are available for download at: <https://osf.io/8ctjy/>.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.ijresmar.2024.09.002>.

References

- Acar, O. A., Gvirtz, A. (2024). *GenAI can help small companies level the playing field*. URL: <https://hbr.org/2024/02/genai-can-help-small-companies-level-the-playing-field>.
- Ammanath, B., Dutt, D., Perricos, C., Sniderman, B. (2024). *Now decides next: Insights from the leading edge of generative AI adoption: Deloitte's State of Generative AI in the Enterprise Quarter one report*. URL: <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/consulting/us-state-of-genai-report.pdf>.
- Andrews, R. L., Ainslie, A., & Currim, I. S. (2002). An empirical comparison of logit choice models with discrete versus continuous representations of heterogeneity. *Journal of Marketing Research*, 39, 479–487. <https://doi.org/10.1509/jmkr.39.4.479.19124>.
- Beichert, M., Bayerl, A., Goldenberg, J., & Lanz, A. (2024). Revenue generation through influencer marketing. *Journal of Marketing*. <https://doi.org/10.1177/00222429231217471>.
- Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Zhuang, J., Lee, J., Guo, Y., Manassra, W., Dhariwal, P., Chu, C., Jiao, Y., Ramesh, A. (2023). *Improving image generation with better captions*. URL: <https://cdn.openai.com/papers/dall-e-3.pdf>.
- Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J.Q., Demszky, D.,...Liang, P. (2021). *On the opportunities and risks of foundation models*. <https://doi.org/10.48550/arXiv.2108.07258>.
- Bonaccorsi, A., & Rossi, C. (2003). Why Open Source software can succeed. *Research Policy*, 32, 1243–1258. [https://doi.org/10.1016/S0048-7333\(03\)00051-9](https://doi.org/10.1016/S0048-7333(03)00051-9).
- Borji, A. (2023). Qualitative failures of image generation models and their application in detecting deepfakes. *Image and Vision Computing*, 137, 104771. <https://doi.org/10.1016/j.imavis.2023.104771>.
- Brynjolfsson, E., Li, D., & Raymond, L. (2023). *Generative AI at Work*. <https://doi.org/10.3386/w31161>.
- Burnap, A., Hauser, J. R., & Timoshenko, A. (2023). Product Aesthetic Design: A Machine Learning Augmentation. *Marketing Science*, 42, 1029–1056. <https://doi.org/10.1287/mksc.2022.1429>.
- Carlson, K., Kopalle, P. K., Riddell, A., Rockmore, D., & Vana, P. (2023). Complementing human effort in online reviews: A deep learning approach to automatic content generation and review synthesis. *International Journal of Research in Marketing*, 40, 54–74. <https://doi.org/10.1016/j.ijresmar.2022.02.004>.
- Ceylan, G., Diehl, K., & Wood, W. (2024). From mentally doing to actually doing: A meta-analysis of induced positive consumption simulations. *Journal of Marketing*, 88, 21–39. <https://doi.org/10.1177/00222429231181071>.
- Chayka, K. (2023). *The Uncanny Failures of A.I.-Generated Hands: When it comes to one of humanity's most important features, machines can grasp small patterns but not the unifying whole*. URL: <https://www.newyorker.com/culture/rabbit-holes/the-uncanny-failures-of-ai-generated-hands>.

- Childers, T. L., & Houston, M. J. (1984). Conditions for a picture-superiority effect on consumer memory. *Journal of Consumer Research*, 11, 643. <https://doi.org/10.1086/209001>.
- Cho, H., Shen, L., & Wilson, K. (2014). Perceived realism: Dimensions and roles in narrative persuasion. *Communication Research*, 41, 828–851. <https://doi.org/10.1177/0093650212450585>.
- Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., Zimmel, R. (2023). *The economic potential of generative AI: The next productivity frontier*. URL: <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>.
- Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4573321>.
- Ding, Y., Tian, C., Ding, H., & Liu, L. (2023). *The CLIP model is secretly an image-to-prompt converter*. <https://doi.org/10.48550/arXiv.2305.12716>.
- Dischler, J. (2023). *Introducing a new era of AI-powered ads with Google*. URL: <https://blog.google/products/ads-commerce/ai-powered-ads-google-marketing-live/>.
- Dzyabura, D., El Kihal, S., Hauser, J. R., & Ibragimov, M. (2023). Leveraging the power of images in managing product return rates. *Marketing Science*, 42, 1125–1142. <https://doi.org/10.1287/mksc.2023.1451>.
- Dzyabura, D., El Kihal, S., & Peres, R. (2022). Image analytics in marketing. In C. Homburg, M. Klarmann, & A. Vomberg (Eds.), *Handbook of market research* (pp. 665–692). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-57413-4_38.
- Dzyabura, D., & Peres, R. (2021). Visual elicitation of brand perception. *Journal of Marketing*, 85, 44–66. <https://doi.org/10.1177/0022242921996661>.
- Feeney, K. (2024). *GenAI ad maker: A new way to create high-quality ads effortlessly*. URL: <https://blog.taboola.com/create-ads-effortlessly/>.
- Feng, X., Zhang, S., & Srinivasan, K. (2023). Marketing through the machine's eyes: Image analytics and interpretability. In K. Sudhir & O. Toubia (Eds.), *Artificial intelligence in marketing* (pp. 217–237). Emerald Publishing Limited. Review of Marketing Research. <https://doi.org/10.1108/S1548-643520230000020013>.
- Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66, 111–126. <https://doi.org/10.1007/s12599-023-00834-7>.
- Gartner (2024). *Gartner experts answer the top generative AI questions for your enterprise: Generative AI isn't just a technology or a business case – It is a key part of a society in which people and machines work together*. URL: <https://www.gartner.com/en/topics/generative-ai>.
- Grewal, R., Gupta, S., & Hamilton, R. (2021). Marketing insights from multimedia data: Text, image, audio, and video. *Journal of Marketing Research*, 58, 1025–1033. <https://doi.org/10.1177/002224372111054601>.
- Hartmann, J., Heitmann, M., Schamp, C., & Netzer, O. (2021). The power of brand selfies. *Journal of Marketing Research*, 58, 1159–1177. <https://doi.org/10.1177/00222437211037258>.
- Hartmann, J., Heitmann, M., Siebert, C., & Schamp, C. (2023). More than a feeling: Accuracy and application of sentiment analysis. *International Journal of Research in Marketing*, 40, 75–87. <https://doi.org/10.1016/j.ijresmar.2022.05.005>.
- Hartmann, J., Huppertz, J., Schamp, C., & Heitmann, M. (2019). Comparing automated text classification methods. *International Journal of Research in Marketing*, 36, 20–38. <https://doi.org/10.1016/j.ijresmar.2018.09.009>.
- Hauser, J. R., Urban, G. L., Liberali, G., & Braun, M. (2009). Website morphing. *Marketing Science*, 28, 202–223. <https://doi.org/10.1287/mksc.1080.0459>.
- Horton, C. B., White, M. W., & Iyengar, S. S. (2023). Bias against AI art can enhance perceptions of human creativity. *Scientific Reports*, 13, 19001. <https://doi.org/10.1038/s41598-023-45202-3>.
- Hui, X., Reshef, O., & Zhou, L. (2023). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4527336>.
- Hulland, J., & Houston, M. (2021). The importance of behavioral outcomes. *Journal of the Academy of Marketing Science*, 49, 437–440. <https://doi.org/10.1007/s11747-020-00764-w>.
- Inman, J. J., Meyer, R. J., Schweidel, D. A., & Srinivasan, R. (2024). Do great powers come with great responsibility? Opportunities and tensions of new technologies in marketing. *International Journal of Research in Marketing*, 41, 18–23. <https://doi.org/10.1016/j.ijresmar.2024.01.006>.
- Jakesch, M., Hancock, J. T., & Naaman, M. (2023). Human heuristics for AI-generated language are flawed. *PNAS*, 120. <https://doi.org/10.1073/pnas.2208839120>. e2208839120.
- Jansen, T., Heitmann, M., Reisenbichler, M., & Schweidel, D. A. (2024). Automated alignment: Engaging customers with visual generative AI. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4656622>.
- Karpinska-Krakowiak, M., & Eisend, M. (2024). Realistic portrayals of untrue information: The effects of deepfaked ads and different types of disclosures. *Journal of Advertising*, 1–11. <https://doi.org/10.1080/00913367.2024.2306415>.
- Keller, K. L. (1993). Conceptualizing, measuring, and managing customer-based brand equity. *Journal of Marketing*, 57, 1–22. <https://doi.org/10.1177/002224299305700101>.
- Keller, K. L., & Lehmann, D. R. (2006). Brands and branding: Research findings and future priorities. *Marketing Science*, 25, 740–759. <https://doi.org/10.1287/mksc.1050.0153>.
- Kelly, C. (2023). *Coke asks consumers to generate art with new AI platform*. URL: <https://www.marketingdive.com/news/coca-cola-coke-generative-ai-marketing-art/645465/>.
- Kim, B. K., Choi, J., & Waksalak, C. J. (2019). The image realism effect: The effect of unrealistic product images in advertising. *Journal of Advertising*, 48, 251–270. <https://doi.org/10.1080/00913367.2019.1597787>.
- King, A. (2024). *AI to the Rescue? – BMG Says a 'Single Project' can involve up to 700 digital assets*. URL: <https://www.digitalmusicnews.com/2024/01/30/bmg-digital-assets-management-ai-to-the-rescue/>.
- Krugmann, J. O., & Hartmann, J. (2024). Sentiment analysis in the age of generative AI. *Customer Needs and Solutions*, 11. <https://doi.org/10.1007/s40547-024-00143-4>.
- Kyriakidi, M. (2022). *Modern marketing dilemmas: Where does performance marketing meet brand building?* URL: <https://www.kantar.com/inspiration/brands/modern-marketing-dilemmas-where-does-performance-marketing-meet-brand-building>.
- Li, Y., & Xie, Y. (2020). Is a picture worth a thousand words? An empirical study of image content and social media engagement. *Journal of Marketing Research*, 57, 1–19. <https://doi.org/10.1177/0022243719881113>.
- Liu, L., Dzyabura, D., Mizik, N. (2020). Visual listening. In: *Extracting brand image portrayed on social media*. *Marketing Science* 39 (pp. 669–686). <https://doi.org/10.1287/mksc.2020.1226>.
- MacKenzie, S. B., Lutz, R. J., & Belch, G. E. (1986). The role of attitude toward the ad as a mediator of advertising effectiveness: A test of competing explanations. *Journal of Marketing Research*, 23, 130–143. <https://doi.org/10.2307/3151660>.
- Matz, S. C., Kosinski, M., Nave, G., & Stillwell, D. J. (2017). Psychological targeting as an effective approach to digital mass persuasion. *PNAS*, 114, 12714–12719. <https://doi.org/10.1073/pnas.1710966114>.
- Matz, S. C., Teeny, J. D., Vaid, S. S., Peters, H., Harari, G. M., & Cerf, M. (2024). The potential of generative AI for personalized persuasion at scale. *Scientific Reports*, 14, 4692. <https://doi.org/10.1038/s41598-024-53755-0>.
- Miller, E. J., Steward, B. A., Witkower, Z., Sutherland, C. A. M., Krumhuber, E. G., & Dawel, A. (2023). AI hyperrealism: Why AI faces are perceived as more real than human ones. *Psychological Science*, 34, 1390–1403. <https://doi.org/10.1177/09567976231207095>.
- Nightingale, S. J., & Farid, H. (2022). AI-synthesized faces are indistinguishable from real faces and more trustworthy. *PNAS*, 119. <https://doi.org/10.1073/pnas.2120481119>. e2120481119.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381, 187–192. <https://doi.org/10.1126/science.adh2586>.

- Oatawa (2024). *User portfolio*. URL: <https://www.istockphoto.com/de/portfolio/oatawa>.
- Paivio, A., & Csapo, K. (1973). Picture superiority in free recall: Imagery or dual coding? *Cognitive Psychology*, 5, 176–206. [https://doi.org/10.1016/0010-0285\(73\)90032-7](https://doi.org/10.1016/0010-0285(73)90032-7).
- Peng, S., Kalliamvakou, E., Cihon, P., & Demirel, M. (2023). *The impact of AI on developer productivity: Evidence from GitHub Copilot*. <https://doi.org/10.48550/arXiv.2302.06590>.
- Peres, R., Schreier, M., Schweidel, D., & Sorescu, A. (2023). On ChatGPT and beyond: How generative artificial intelligence may affect research, teaching, and practice. *International Journal of Research in Marketing*, 40, 269–275. <https://doi.org/10.1016/j.ijresmar.2023.03.001>.
- Phillips, B. J., McQuarrie, E. F., & Griffin, W. G. (2014). How visual brand identity shapes consumer response. *Psychology & Marketing*, 31, 225–236. <https://doi.org/10.1002/mar.20689>.
- Pieters, R., & Wedel, M. (2004). Attention capture and transfer in advertising: Brand, pictorial, and text-size effects. *Journal of Marketing*, 68, 36–50. <https://doi.org/10.1509/jmkg.68.2.36.27794>.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., & Sutskever, I. (2021). Learning transferable visual models from natural language supervision. *Proceedings of Machine Learning Research*, 8748–8763 <https://proceedings.mlr.press/v139/radford21a>.
- Reisenbichler, M., Reutterer, T., Schweidel, D. (2023). *Applying large language models to sponsored search advertising*. URL: <https://www.msi.org/working-paper/applying-large-language-models-to-sponsored-search-advertising/>.
- Reisenbichler, M., Reutterer, T., Schweidel, D. A., & Dan, D. (2022). Frontiers: Supporting content marketing with natural language generation. *Marketing Science*, 41, 441–452. <https://doi.org/10.1287/mksc.2022.1354>.
- Rizzo, G. L. C., Berger, J. A., & Villarreal Ordenes, F. (2023). What drives virtual influencer's impact? *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4329150>.
- Rodgers, B. (2021). *How much does commercial product photography cost?* URL: <https://digitalartthatrocks.com/blog/2021/12/8/how-much-does-commercial-product-photography-cost>.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In: *2022 IEEE/CVF Conference on computer vision and pattern recognition (CVPR)*, IEEE (pp. 10674–10685). <https://doi.org/10.1109/CVPR52688.2022.01042>.
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M. (2022). Photorealistic text-to-image diffusion models with deep language understanding. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Eds.), *Advances in neural information processing systems*. URL: https://proceedings.neurips.cc/paper_files/paper/2022/hash/ec795aeadae0b7d230fa35cbaf04c041-Abstract-Conference.html.
- Sauer, A., Lorenz, D., Blattmann, A., & Rombach, R. (2023). *Adversarial diffusion distillation*. <https://doi.org/10.48550/arXiv.2311.17042>.
- Schwartz, E. M., Bradlow, E. T., & Fader, P. S. (2017). Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science*, 36, 500–522. <https://doi.org/10.1287/mksc.2016.1023>.
- Seiler, S., Tuchman, A., & Yao, S. (2021). The impact of soda taxes: Pass-through, tax avoidance, and nutritional effects. *Journal of Marketing Research*, 58, 22–49. <https://doi.org/10.1177/0022243720969401>.
- Sloane, G. (2024). *Brands add AI restrictions to agency contracts—Behind the growing trend*. URL: <https://adage.com/article/digital-marketing-ad-tech-news/ai-restrictions-added-ad-agency-contracts/2548696>.
- Smith, R. E., MacKenzie, S. B., Yang, X., Buchholz, L. M., & Darley, W. K. (2007). Modeling the determinants and effects of creativity in advertising. *Marketing Science*, 26, 819–833 <http://www.jstor.org/stable/40057228>.
- Steenkamp, J. B. E., van Heerde, H. J., & Geyskens, I. (2010). What makes consumers willing to pay a price premium for national brands over private labels? *Journal of Marketing Research*, 47, 1011–1024. <https://doi.org/10.1509/jmkr.47.6.1011>.
- Talebi, H., & Milanfar, P. (2018). NIMA: Neural image assessment. *IEEE Transactions on Image Processing*. <https://doi.org/10.1109/TIP.2018.2831899>.
- The One Club (2023). *AI Ketchup*. URL: <https://www.oneclub.org/awards/theoneshow/-award/48285/ai-ketchup/ai-ketchup>.
- Thomas, D. (2024). *AI advertising start-up valued at \$4bn after fundraising*. URL: <https://www.ft.com/content/4c7bee10-51d3-489b-873a-765157af8aac>.
- van Heerde, H. J., Moorman, C., Moreau, C. P., & Palmatier, R. W. (2021). Reality check: Infusing ecological value into academic marketing research. *Journal of Marketing*, 85, 1–13. <https://doi.org/10.1177/0022242921992383>.
- Wang, W., Bell, J. J., Dotson, J. P., & Schweidel, D. A. (2023). Generative AI and artists: Consumer preferences for style and fair compensation. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4428509>.
- Xu, L., & Mehta, R. (2022). Technology devalues luxury? Exploring consumer responses to AI-designed luxury products. *Journal of the Academy of Marketing Science*, 50, 1135–1152. <https://doi.org/10.1007/s11747-022-00854-x>.
- Zhang, H., Bai, X., & Ma, Z. (2022a). Consumer reactions to AI design: Exploring consumer willingness to pay for AI-designed products. *Psychology & Marketing*, 39, 2171–2183. <https://doi.org/10.1002/mar.21721>.
- Zhang, M., & Luo, L. (2023). Can consumer-posted photos serve as a leading indicator of restaurant survival? Evidence from Yelp. *Management Science*, 69, 25–50. <https://doi.org/10.1287/mnsc.2022.4359>.
- Zhang, S., Lee, D., Singh, P. V., & Srinivasan, K. (2022b). What makes a good image? Airbnb demand analytics leveraging interpretable image features. *Management Science*, 68, 5644–5666. <https://doi.org/10.1287/mnsc.2021.4175>.
- Zhang, X., Zhou, M., & Lee, G. M. (2024). Generative AI and creator economy: Investigating the effects of AI-generated voice on online video creation. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4676705>.
- Zhou, E., & Lee, D. (2024). Generative artificial intelligence, human creativity, and art. *PNAS Nexus*, 3, pgae052. <https://doi.org/10.1093/pnasnexus/pgae052>.