JACOBS UNIVERSITY

MSc. Data Science for Business and Society

MDSSB-APP-01 Current Topics and Applications in Data Science

Final Term Paper

The Generative AI Revolution and its Impact on Marketing

Course Instructors: Prof. Dr. Hilke Brockmann and Prof. Dr. Adalbert Wilhelm

Handed in by: Suraj Karakulath

Date: 31.12.2022

Table of Contents

Abstract **3**

Background: AI in a marketing context **3**

Definition: What is generative AI? **4**

The evolution of generative models **5**
- Generative Adversarial Neural Networks (GANs) 5
- Transformers 6
- Generative Pre-trained Transformer (GPT)-3 6

Applications of generative AI in marketing **6**
- Text generation for campaigns and website content 7
- Creative social media campaigns 7
- Scaling video production 9
- Personalised customer support through chats 9

Challenges and limitations of generative AI **10**
- Explainability and the “black box” conundrum 10
- Errors due to training data 11
- Bias and discrimination 11

A threat to the creative workforce? **11**

Conclusion **12**

Abstract

There is no doubting the fact that automation and artificial intelligence (AI) has disrupted many industries which relied on routine, repetitive tasks, including the marketing sector. But a number of recent research papers and innovations in the area of generative AI are creating much excitement along with concern, signalling a major paradigm shift in the creative aspect of marketing, an area which was thought to be off-limits to disruption from AI, at least for a while. In this paper, we provide a brief overview of how AI has so far been used in the marketing sector and how the rise of generative AI will transform the industry. We define generative AI and track its evolution as well as some of the major milestones over the last decade that have accelerated its development. We then discuss the various ways in which generative AI is being used in the marketing industry, including examples of successful applications, which showcase some of the benefits that businesses and marketing agencies will reap from this technology. We also touch upon some of the challenges and limitations of generative AI in the marketing industry, including issues related to explainability and bias, as well as the threat it poses to the creative workforce.

Background: AI in a marketing context

Marketing as a discipline has for very long been seen as both an art and a science. With the goal of reaching a wide audience with relevant messaging, marketing relies on creativity to appeal to people’s emotions for decision-making around purchase and experimentation to analyse what type of creative content performs the best in achieving this goal.

Many technologies over the decade have helped in the scientific aspect of marketing, such as digital platforms that collect first-party data on the behaviour of visitors to a website and ad platforms that streamline and automate the delivery of ads to targeted audiences. During the Mad Men era of the 1960s, marketing was mainly about advertising campaigns driven by the creative visions of strategists, copywriters and artists. This changed in the last few decades with digital marketing and platforms like Google Ads and Facebook. Instead of several months spent on advertising campaigns on traditional media such as print, TV and radio, marketing became easier and cheaper to test and experiment with digital ads online. It

also reduced the timeframe for a campaign from around a year or more to a mere few weeks.

But even in the digital marketing era, the creative aspect was still important, as it was still necessary for artists, designers and writers to come up with novel ideas for the ads before testing them. This is slowly changing with the arrival of generative AI and the ease of access to various tools that automate content creation. In fact, Gartner, a renowned technological research and consulting firm included generative AI in their recently released list of seven technology disruptions that will impact sales through 2027[1].

Definition: What is generative AI?

Generative AI is a subfield of AI that focuses on generating new data or content based on learned patterns and distributions. According to Gartner's description, generative AI "learns from existing content artefacts to generate new, realistic artefacts that reflect the characteristics of the training data, but do not repeat it." The content produced by generative AI can be anything from text, images, video, music, speech or even software code. This is done by using complex machine learning models to predict the next word based on previous word sequences, or the next image based on words describing previous images.

In the marketing sector, generative AI can broadly be categorised into three types.

- **Text-to-text:** These produce text as output based on text prompts. Think of blog posts, social media status and other marketing copy. The most recent ChatGPT tool from OpenAI is an example of such a generative AI, which aims to produce text in a conversational format based on prompts entered by humans in natural language.
- **Text-to-image:** These produce images, either photorealistic or illustrative or other creative styles based on text prompts. DALL-E, another tool from OpenAI is an example of this type of generative AI.
- **Text-to-video:** These produce videos from text, with Meta's Make-A-Video being one of the most recent examples.

While this is not a new invention or innovation, what is striking is how each successive development in the technology over the last few years and decade has improved over the previous version.

The evolution of generative models

The history of generative AI is indeed closely tied to that of AI, dating back to the early days of research in the field. Some of the earliest efforts in the area of generative AI were seen in statistical techniques such as n-gram models, which estimate the probability of a word based on the probability of the previous n-1 words in a sequence. These models were limited in their ability to capture long-range dependencies and had poor performance on tasks requiring a deep understanding of language.

In the late 1980s and early 1990s, the introduction of neural networks and deep learning techniques sparked a resurgence of interest in generative AI. These techniques, which are inspired by the structure and function of the brain, enabled the creation of more powerful generative models that could learn to generate data based on large amounts of training data. One early example of a generative neural network is the Restricted Boltzmann Machine (RBM), introduced by Hinton et. al in 2006[2], which can learn to generate synthetic data that is similar to a given dataset.

Generative Adversarial Neural Networks (GANs)

In the 2010s, there were many more developments in the evolution of generative AI but two were significant in how they accelerated the research. The first was the introduction of Generative Adversarial Neural networks or GANs by Ian Goodfellow et al. in 2014, providing a powerful new approach for generating synthetic data that is similar to real-world data.

GANs consist of two neural networks[3]: a generator network that generates synthetic data, and a discriminator network that determines whether a given sample is real or synthetic. The two networks are trained together in an "adversarial" manner, with the generator trying to generate synthetic data that is indistinguishable from real data, and the discriminator trying to distinguish between the two. This process leads to the generator network learning to generate high-quality synthetic data that is similar to the real data.

Transformers

Another major development was the introduction of transformative language models, such as the Transformer model introduced by Vaswani et al. in 2017[4], began to emerge. This was a revolutionary new approach to natural language processing, based on the idea of

self-attention, which allows the model to directly model relationships between input elements, rather than relying on sequential processing. This made the Transformer model particularly well-suited for tasks that require a deep understanding of language, such as translation and language modelling.

Generative Pre-trained Transformer (GPT)-3

Following this, language models continued to get better at performing well on a wide range of tasks. Then in 2020, the GPT-3 model was introduced by Brown et al.[5] which was trained on 45 terabytes of data and employed 175 billion parameters or coefficients to make its predictions. This model was capable of producing very impressive results on a variety of tasks, including language translation, text generation, and question answering. GPT-3 significantly improved the state of the art in natural language processing.

Applications of generative AI in marketing

Technology has been encroaching on various aspects of marketing over the years. But these have mostly been restricted to the use of data science for generating insights from large datasets of customer or audience data or machine learning models for the prediction of sales forecasts or at the very best the automation of repetitive tasks such as optimisation of ads. Generative AI however, is a game-changer as it is poised to revolutionise the field of marketing by enabling the creation of content for various different marketing operations.

Text generation for campaigns and website content

Generative AI can be used to create text-based content for marketing campaigns. There is a rising demand for personalised digital content today and generative AI can help businesses and organisations scale their content creation without compromising on quality. According to InsightSLICE, the global digital creation market is expected to grow 12% annually between 2020 and 2030 and hit $38.2 billion[6]. This rise in demand means that traditional ways that marketing is done may not be fast enough, and this is one area where generative AI can be helpful. For example, Jasper is a generative AI company built on GPT-3 that focuses on producing blogs, website copy, marketing emails and ad copy among other types of content. It also claims to be able to produce content that is optimised for search engines[7].

Creative social media campaigns

Along with text-to-text models, text-to-image models are also being used to produce visual creative assets for marketing campaigns. Heinz, for example, used the image generator DALL-E 2 for the first ever ad campaign that used visuals generated by AI[8]. Using prompts such as “ketchup in outer space” and others sourced from fans of the brand, it created a series of posts with an image of a ketchup bottle and a label similar to that of the Heinz logo with the caption “This is what ‘ketchup’ looks like to AI.”

Source: *Heinz A.I. Ketchup*, *Campaigns of the World*, August 2022

Nestle on the other hand partnered with the French brand La Laitière and ad agency Ogilvy to use the DALL.E outpainting feature for a creative campaign titled “It’s so pleasurable to take the time”[9]. This campaign used generative AI to create paintings in the style of the Dutch artist Johannes Vermeer.

Source: *It's so pleasurable to take the time, Ads of the World,* September 2022.

Scaling video production

More audiences today are consuming and indeed, preferring video content, making it imperative for marketing campaigns to be producing videos. And yet, video production has always been time- and resource-intensive parts of campaigns. The entire process from pre-production to location scouting, scheduling, videography and editing takes countless hours which also drives up the costs.

Generative AI can be used to make this process more efficient and scale up the volume of production. Today's state-of-the-art text-to-video AI models can help create videos that look real and which are also relatively cheap. Moreover, they can also be personalised so that in an attempt to reach wider audiences, the relevancy of the content is not lost. For example, one campaign from PepsiCo's involved a deep fake of Lionel Messi, in which he appears to send personalised videos to fans, and in different languages[10].

Source: messimessages.com, UEFA Champions League in partnership with Lays

Personalised customer support through chats

Generative AI can be used to automate certain tasks, such as answering frequently asked questions or generating responses to customer inquiries. For example, a generative AI system could be trained on a dataset of customer inquiries and responses, and then be used to generate responses to new customer inquiries in real-time. This can improve the efficiency and accuracy of customer support, and free up human customer service agents to focus on more complex tasks. Some organisations like Bank of America, are already doing this. Their virtual assistant is powered by generative AI and can communicate with customers via text, voice, or by tapping options in a natural language conversation.

Challenges and limitations of generative AI

Even as generative AI opens up these numerous applications, its use is still not a silver bullet to all problems in marketing and it is certainly not without limitations. In this section, we will explore these challenges and limitations and discuss their implications for the use of generative AI in marketing. Understanding these challenges and limitations is essential for ensuring that the benefits of generative AI are balanced with the needs of consumers and the creative workforce.

Explainability and the “black box” conundrum

One of the most significant concerns with AI is the problem of explainability. Explainability refers to the ability to understand and interpret the processes and decisions made by a generative AI model. This is important in general because it allows stakeholders, such as end users, regulatory bodies, and developers, to understand exactly how the model is making predictions or generating data so that were there any potential issues or biases, it can still be explored and resolved.

However, this is not an easy task, as many AI models, including generative models, are "black boxes" that are difficult to interpret and understand. Consider a text-to-image generative AI tool for example. Once a text prompt is typed in as input, a few images that fit the prompt are generated. If the text is repeatedly modified further, the model may produce something very close to what was intended.

But even if the model produces the exact image that was expected, there is no way to know or understand how it is generated. All that is known is that it involves neural networks, and in the case of GANs, one trained to match an image with a descriptive text and another to generate images from scratch. Once the second neural network creates an image that the first one accepts as a match for the prompt, the results are satisfactory. For example, consider the text prompt “a panda wearing a cowboy hat in a garden”. We may get exactly what that describes but we don’t know what characteristics, attributes or pixels in the images are sufficient for the model to “know” what a panda is, what a “cowboy hat” is and what a garden is.

Errors due to training data

There are the occasional errors as well in the results of generative AI models. For example, when a prompt “salmon in a river” was entered into one of the text-to-image tools, instead of salmon fish swimming in a river, the image showed a strange collection of chopped-up fillets floating downstream. Due to the same black box conundrum, generative models may not understand the meaning of what they produce, no matter how accurate the results look. According to the renowned linguist Noam Chomsky, a generative model like GPT-3 does not produce words in a meaningful language any differently from how it would produce words in a meaningless or impossible language.

Bias and discrimination

Yet another issue that marketers may need to be wary of is the problem of bias. As is the case with most AI models, generative AI may also be biased or discriminatory if they are trained on biased or discriminatory data. For instance, a model trained on data that is skewed towards a specific demographic or race may generate biased or discriminatory content. And if content generated by such models is used for mass-outreach ad campaigns, it can backfire and result in negative consequences for both the business and its customers.

A threat to the creative workforce?

As with any radical new technology, and especially one in the area of AI and automation, generative models are raising alarms among the creative workforce all over the world. While there has been a concern about AI making many jobs redundant, it has mainly been jobs which involved routine, repetitive tasks. Any creative work, or jobs that involved generating new ideas, content or visuals were safe for the foreseeable future, or at least that was the perception until the most recent developments in generative AI.

In fact, artists have already raised their concerns and even taken to legal recourse in the wake of the introduction of DALL-E 2 in April 2022 and additional text-to-video models from Google and Meta in October 2022[11]. In what they see as “one of the biggest upheavals in a generation”, they fear that some of them will lose work. Some artists are already fighting legal battles over what they consider to be misappropriation of images to train models that could replace them.

Even stock image licensing firms are taking action to preserve their competitive advantage. If generative models can create images based on any prompts provided by users, then the very idea of stock images that only show predefined scenarios will become redundant altogether. Moreover, there are also copyright concerns, as it is reasonable to expect that generative AI models are capable of scraping publicly available photos for the purpose of training their algorithms before generating new imagery. These publicly available images may also include copyrighted ones that artists and photographers submit to stock photo websites. And if the newly generated images are used, the question of who gets to own these and the copyright issues surrounding them fall in an extremely grey area.

Getty Images, one of the leading stock image providers, announced in September 2022 that it would ban AI-generated art, including those produced by OpenAI's DALL-E and Meta AI's Make-A-Scene, from its platform[12]. Shutterstock, another stock image provider, is taking a different approach and deciding to partner with generative AI, as it signed a deal with OpenAI in October 2022 to integrate DALL-E 2 directly into their platform[13]. They also plan to launch a "Contributor Fund" to reimburse creators when their work is sold to train text-to-image AI models.

Conclusion

Technology has always pushed the boundaries of what humankind can do. With every new invention and innovation, right from the development of steam engines all the way to the widespread adoption of the internet and the democratisation of computing, there has been a mix of excitement for the possibilities and concern for the negative consequences. Generative AI is no different, as it opens the doors to numerous possibilities for creators. In the field of marketing, generative AI will be game-changing for many applications, including content generation, social media management, and personalised marketing. At the same time, it will also raise concerns about explainability, bias, copyright and ethical issues.

Businesses will no doubt experiment with generative AI tools in search of efficiency, productivity and cost savings. It is highly likely that there will be a large cohort of creators who will be replaced, especially those who work on the repetitive tasks in marketing such as product description, pixel pushing in design and templatised video creation and editing. It is important for creators to stay one step ahead and adapt to their changing role in marketing and explore how they can work together with AI tools. This could involve upskilling on the use of such tools, studying the results of various models to improve their "prompt engineering" skills and finding ways to make themselves more valuable than the mere output of a generative AI model.

References

1. https://www.gartner.com/en/newsroom/press-releases/2022-10-10-gartner-identifies-seven-technology-disruptions-that-will-impact-sales-through-2027
2. https://www.cs.toronto.edu/~hinton/absps/ncfast.pdf
3. https://arxiv.org/abs/1406.2661

4. https://arxiv.org/abs/1706.03762
5. https://arxiv.org/abs/2005.14165
6. https://www.insightslice.com/digital-content-creation-market
7. https://hbr.org/2022/11/how-generative-ai-is-changing-creative-work
8. https://campaignsoftheworld.com/digital/heinz-a-i-ketchup/
9. https://www.adsoftheworld.com/campaigns/it-s-so-pleasurable-to-take-the-time
10. https://www.sportsbusinessjournal.com/Daily/Issues/2021/03/19/Marketing-and-Sponsorship/Messi-Lays.aspx
11. https://www.technologyreview.com/2022/12/16/1065005/generative-ai-revolution-art/
12. https://www.artnews.com/art-news/news/getty-images-bans-ai-generated-images-due-to-copyright-1234640201/
13. https://www.theverge.com/2022/10/25/23422359/shutterstock-ai-generated-art-openai-dall-e-partnership-contributors-fund-reimbursement

Citations

Gartner Identifies 7 Technology Disruptions That Will Impact Sales Through 2027. (2022, October 10). Gartner. Retrieved December 29, 2022, from https://www.gartner.com/en/newsroom/press-releases/2022-10-10-gartner-identifies-seven-technology-disruptions-that-will-impact-sales-through-2027

Hinton, G. E., Osindero, S., & Teh, Y.-W. (n.d.). A Fast Learning Algorithm for Deep Belief Nets. *Neural Computation*, *18*(7), 1527-1554. Retrieved December 27, 2022, from https://www.cs.toronto.edu/~hinton/absps/ncfast.pdf

1406.2661] Generative Adversarial Networks. (2014, June 10). arXiv. Retrieved December 21, 2022, from https://arxiv.org/abs/1406.2661

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., Polosukhin, I. Attention Is All You Need. *Advances in Neural Information Processing Systems, 30*. Retrieved December 21, 2022, from https://arxiv.org/abs/1706.03762

Brown, T. B. (2020, May 28). *[2005.14165] Language Models are Few-Shot Learners*. arXiv. Retrieved December 27, 2022, from https://arxiv.org/abs/2005.14165

Digital Content Creation Market to Grow at 12% CAGR by 2030. (n.d.). insightSLICE. Retrieved December 16, 2022, from https://www.insightslice.com/digital-content-creation-market

How Generative AI Is Changing Creative Work. (2022, November 14). Harvard Business Review. Retrieved December 16, 2022, from https://hbr.org/2022/11/how-generative-ai-is-changing-creative-work

Heinz A.I. Ketchup | Just like humans, A.I. prefers Heinz. (2022, August 17). Campaigns of the World®. Retrieved December 16, 2022, from https://campaignsoftheworld.com/digital/heinz-a-i-ketchup/

Nestlé: It's so pleasurable to take the time • Ads of the World™ | Part of The Clio Network. (2022, September 20). Ads of the World. Retrieved December 16, 2022, from https://www.adsoftheworld.com/campaigns/it-s-so-pleasurable-to-take-the-time

Lay's app uses deepfake tech for custom messages from Lionel Messi. Sports Business Journal. Retrieved December 16, 2022, from https://www.sportsbusinessjournal.com/Daily/Issues/2021/03/19/Marketing-and-Sponsorship/Messi-Lays.aspx

Douglas, W. (2022, December 16). *Generative AI is changing everything. But what's left when the hype is gone?* MIT Technology Review. Retrieved December 16, 2022, from https://www.technologyreview.com/2022/12/16/1065005/generative-ai-revolution-art/

Jacobs, H. (2022, September 22). *Getty Images Bans AI-Generated Images Due To Copyright Worries*. Art News. Retrieved December 16, 2022, from https://www.artnews.com/art-news/news/getty-images-bans-ai-generated-images-due-to-copyright-1234640201/

Vincent, J. (2022, October 25). *Shutterstock will start selling AI-generated stock imagery with help from OpenAI*. The Verge. Retrieved December 16, 2022, from https://www.theverge.com/2022/10/25/23422359/shutterstock-ai-generated-art-openai-dall-e-partnership-contributors-fund-reimbursement