

# Moving Beyond ChatGPT: Applying Large Language Models in Marketing Contexts

## THE AUTHORS

### David A. Schweidel

Rebecca Cheney McGreevy Endowed Chair and Professor of Marketing, Goizueta Business School, Emory University

### Martin Reisenbichler

Postdoctoral Researcher, University of Hamburg,


### Thomas Reutterer

Professor of Marketing, Vienna University of Economics and Business

## KEYWORDS

Large Language Models, Content Marketing, Search Engine Optimization, Human-Machine Collaboration





*AI is a tool that can  
boost productivity and  
performance, but it requires  
human oversight.*



**A new era of AI assistance** ✕ When Open AI demonstrated ChatGPT to the public, people were amazed by what large language models (LLMs) – the type of generative AI behind the chat-like surface – were able to produce. Even if LLMs might seem like sentient machines, they should more appropriately be viewed as “stochastic parrots” or eager-to-please interns. But despite their apparent prowess, LLMs are not trained for a particular context. Should the text attract new customers? Engage current customers? Will it be used for direct mail or a blog post? The intended use of the text will ultimately dictate what makes for successful content. If we could filter previously developed content and use only that which has been deployed successfully for a particular task, we could try to recreate the formula for success. This is not wishful thinking, but rather offers an accessible approach to tailoring LLMs for marketing applications that we have successfully demonstrated in the search engine marketing process.

**LLMs perform well “inside the box”** ✕ Figure 1 offers a schematic of the types of projects that stand to benefit most and least from LLMs. To harness LLMs like OpenAI’s GPT Series or Google’s Gemini, we must reframe “out of the box” problems as “inside the box” concepts where LLMs shine. “Inside the box” thinking makes use of the patterns present in existing data. Ask an LLM for the recipe for a traditional dessert, and it has many examples on which to draw. If you wanted to blend recipes for different confections, LLMs will draw upon relations observed in the training data to craft a unique combination. The same would be true for drafting descriptions of new products that evolved from prior products.

**FIGURE 1** > The Opportunities for LLMs in Marketing




»

*Creating new content by leveraging past content quickly escalates beyond humans' pattern recognition capabilities.*

«

The situation is different for fundamentally new tasks. Imagine trying to describe today's smartphones before they came on the market. The device could have easily been described as a miniature computer with a powerful camera that just happens to have the capability of making phone calls. We still talk today of "dialing" a number, even though many have never seen a rotary phone. While it is easy to deconstruct today's products into underlying components à la conjoint analysis,

put yourself in the early days of telephones. Anyone writing about a phone at the time would have had no concept of the capabilities of today's smartphones. Envisioning new products requires "outside the box" thinking that goes beyond what has been previously observed. Such paradigm shifts reveal the shortcomings of being anchored to historical data. If it has never been seen before, there is no data on which to train an AI. Such creativity hasn't been replaced (yet).



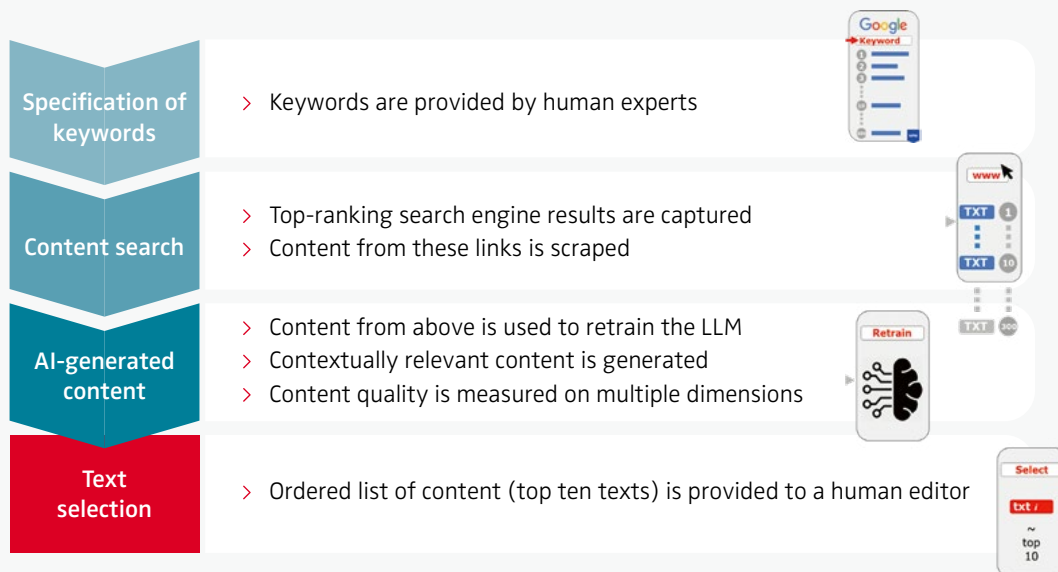
## BOX 1

## Customizing LLMs by adding context-specific information and human editing


Our approach to producing AI-assisted SEO content blends the general language patterns known to a base LLM with application-specific language. The result is a customized language model that can generate thousands of pieces of hypothetical content for a particular search query.

But not all content from our informed “stochastic parrot” will be suitable for publishing. Some might be gibberish, or it might replicate toxic or otherwise inappropriate language, while other pieces of content will be more appealing. To find the text that is expected to perform best, we score the content using an empirically validated quality metric that was developed by scraping nearly 1.5 million websites using 8,500 search terms that spanned 36 different industries. While the content created up until this point is fully automated, it is risky to let machine-generated content represent your brand. Content that is factually accurate – which is not guaranteed – may not be consistent with the tone or positioning of the brand. An experienced editor’s discerning eye is essential to choosing the appropriate content and casting it into a form that is publishable. Figure 2 summarizes the steps in this process.

FIGURE 2 > Overview steps of a (semi)automated SEO content writing machine



**Performance of the hybrid solution** ✕ To evaluate this approach, we pitted our human-in-the-loop AI against SEO experts. In two different industry contexts, our approach outperformed SEO experts. The machine-generated-human-edited content was indistinguishable from human-written content in terms of readability, ease of understanding, naturalness and credibility. More importantly, our content ranked higher in the search engine than human-written content, outperforming it in terms of online visibility by more than a factor of 2. Beyond the improved performance, the machine-generated, human-edited content was created at a fraction of the cost compared to hiring a professional copywriter. Even if we account for the effort to build and maintain the AI solution and the labor costs associated with editing the AI-generated content, the cost savings per piece of content exceeded 80%.



### LLMs excel when human and machine collaborate

✘ While AI can produce candidate answers to a question at scale, human experience allows us to ask the right questions and to inform how we measure success. Take a non-profit trying to increase donations to its fundraising campaign. While some past appeals worked to open constituents' wallets, others fell on deaf ears. We would want to draw inspiration from the marketing content that was successful. By first defining and identifying past successes, we can use that to provide a schematic for AI-assisted content creation. But textual content is multi-faceted as it involves tone, sentiment, word frequency and syntactic patterns. Creating new content by leveraging past content quickly escalates beyond humans' pattern recognition capabilities. This is where LLMs excel, leveraging previously observed patterns to deliver fresh content bearing these same language patterns. The

approach we developed and implemented, as summarized in Box 1, relies on the pattern recognition capabilities of AI and a human's discerning eye.

**Creating content optimized for online search** ✘ In our research, we develop an AI-supported algorithm to produce search engine optimized (SEO) content, a typical highly scalable "inside the box" application in the upper-right quadrant of Figure 1. The objective was simple: Create content that would rank well in search engines. Being highly ranked is paramount for marketers because search engines are one of the ways to attract new prospects. But it is not enough to be in the search results. The vast majority of clicks occur on the first page of search results, making the subsequent pages of considerably less value. Our goal in creating new content was to show up on the first page.



*Beyond the improved performance, the machine-generated, human-edited content was created at a fraction of the cost compared to hiring a professional copywriter.*





*While LLMs enable us to offload a good portion of the work, marketing still requires a human touch.*



One would expect content similar to the current top-performing search results to perform well. But replicating the current search engine results does not provide value to consumers searching for a particular query. In fact, search engines actively discourage this black hat SEO practice, referred to as content spinning. Instead, it is necessary to create unique content that retains the syntactic patterns of top-performing search results.

One option would be to build an LLM entirely from scratch for each application. This approach, however, is cost-prohibitive. Moreover, the content that is most relevant to a particular application – the currently top-performing websites based on their search results – is fairly limited. Another option would be to use an existing, off-the-shelf LLM, but such models don't include options to tailor the language to our specific task or brand. While the content might look natural to a reader, it hasn't been engineered to perform well with search engines. Therefore, we developed a middle-ground approach by dynamically fine-tuning an existing LLM based on a given search query. It is described in Box 1.

**Explaining the performance disparity between SEO experts and our approach** ✕ While humans might look for easy and straightforward rules of thumb like integrating some keywords, machine-generated content recreates high-dimensional textual properties like topical fit and optimal keyword frequency as well as industry-specific lexical and reading patterns. This is the point where a human's analytical gaze reaches its limits and LLMs excel. But it is important to remember that the machine cannot go at it alone. Though we have reduced the effort needed from human content creators, the role remains critical to ensure accuracy and consistency with brand guidelines. As we demonstrate, while LLMs enable us to offload a good portion of the work, marketing still requires a human touch.

**Early adopters are thrilled** ✕ Among our early adopters, the practical benefit is undisputed. For example, Storyblok, a globally operating headless CMS provider, was thrilled by the "overdue arrival of tools for automated marketing content generation." Their evaluation ranged from deeming it "an

interesting source of inspiration" to "extremely useful in conjunction with interactive editing." Christoph Wendl, CEO of Iphos IT solutions and inventor of the enterprise search software searchIT, credited gaining a substantial contract for his company to the superior search engine performance of the content generated as part of our pilot studies. The substantial reduction in costs from automating more routine tasks allows human contributors to focus their efforts on higher-value tasks like tailoring content to a specific brand's narrative or unique proposition.

**Many more applications on the horizon** ✕ Beyond SEO, we have applied similar principles to applications including search engine advertising, display advertising and social media posts. Without a doubt, we will see automated content generation used to optimize entire pieces of creative marketing – text, images, audio and video. But marketers should not be lulled into a "set it and forget it" mindset. Ultimately, they are responsible for the content – whether it is created by human, machine or both. Generative AI is ultimately a tool that can boost productivity and performance, but it has yet to become a wholesale replacement for a marketer's knowledge and skill set. ✕



#### FURTHER READING

Reisenbichler, M., Reutterer, T., Schweidel, D. A., & Dan, D. (2022). Frontiers: Supporting Content Marketing with Natural Language Generation. *Marketing Science*, 41(3), 441–452.