

# LEVERAGING AI FOR CONTENT GENERATION: A CUSTOMER EQUITY PERSPECTIVE

David A. Schweidel, Martin Reisenbichler,  
Thomas Reutterer and Kunpeng Zhang

## ABSTRACT

*Advances in artificial intelligence have ushered in new opportunities for marketers in the domain of content generation. We discuss approaches that have emerged to generate text and image content. Drawing on the customer equity framework, we then discuss the potential applications of automated content generation for customer acquisition, relationship development, and customer retention. We conclude by discussing important considerations that businesses must make prior to adopting automated content generation.*

**Keywords:** Artificial intelligence; language models; natural language generation; generative artificial intelligence; marketing technology; customer equity

Recent advances in artificial intelligence (AI) have given rise to tools that can be used to support marketing content. Chief among these is the development of large natural language generation (NLG) models such as BERT (Bidirectional Encoder Representations from Transformers; [Devlin, Chang, Lee, & Toutanova, 2018](#)), GPT-3 (Generative Pre-trained Transformer 3; [Brown et al., 2020](#)), and Megatron-Turing NLG.<sup>1</sup> Drawing on a massive corpus of digitized text, these models have the potential to generate language for different applications. Following Microsoft's investment in Open AI and the subsequent development of the GPT-3 API, numerous marketing-related businesses have emerged. Product descriptions, search engine advertising, social media posts, and chatbots are just

---

Artificial Intelligence in Marketing

Review of Marketing Research, Volume 20, 125–145

Copyright © 2023 David A. Schweidel, Martin Reisenbichler, Thomas Reutterer and Kunpeng Zhang

Published under exclusive licence by Emerald Publishing Limited

ISSN: 1548-6435/doi:[10.1108/S1548-643520230000020006](https://doi.org/10.1108/S1548-643520230000020006)

some of the use cases for which startups have developed AI-based solutions to create unique content.

But, teaching a machine to write in English (or any language, for that matter) is just the start. Deep learning models have been developed that establish the link between text and images. Drawing on a large library of images that have been associated with text, text prompts can be used to create images (e.g., DALL-E 2, Midjourney, Stable Diffusion, Imagen). Content generation that originates with images has also been applied to develop entirely new images, such as transforming rudimentary sketches to images that resembles pieces created by artists (Nvidia Canvas)<sup>2,3</sup>. Similar techniques are being applied to voices (e.g., Lovo.ai), enabling the development of custom voices and the creation of a market in which participants can license the use of their voices to businesses.

Where do these AI-based tools fit into the marketing landscape? To address this question, we draw upon the customer equity framework (see, e.g., [Gupta & Lehmann, 2005](#); [Rust, Lemon, & Zeithaml, 2004](#)). Viewing the value of the firm as being tied to the value of the customer base, we consider this in terms of the discounted future cash flow of current and prospective customers. Underlying the expected value of an individual customer and the aggregated value across all customers are three specific processes that must be considered: acquisition, relationship development, and retention.

Prospective customers must be attracted before they are ultimately acquired. Marketing efforts may move customers down the purchase funnel, initially building brand awareness and increasing consumer interest. Once acquired, firms take efforts to cultivate the relationship for those selected customers with high projected potential, increasing the depth and breadth of the relationship. The content used to develop a customer relationship in its early stages (e.g., in the context of onboarding new customers) likely differs from marketing communications that loyal consumers will receive once a relationship has been formed. The marketing content employed to strengthen the relationship may also differ from the content that is used to salvage a damaged relationship, to prevent valuable customers from churn or to win back lost customers. That is, the marketing content used to cultivate and enhance an existing relationship may be different from the content used for the purposes of increasing customer retention.

In the next section, we begin by identifying the marketing activities supported by content that are intended to support each component of customer equity. We then discuss the importance of customer profiling and personalization as it applies to acquisition, development, and retention. In taking stock of the content's role within each of these components of customer equity, we then turn our attention to the established and emerging methodologies, as they have been applied to each component. We conclude with a discussion of the opportunities that lie ahead for content generation to support marketing efforts.

## 1. THE POTENTIAL FOR AI THROUGHOUT THE CUSTOMER JOURNEY

Lemon and Verhoef (2016) depict the customer journey as comprising a sequence of experiences that a customer has with a brand, with each experience consisting of three stages (prepurchase, purchase, and postpurchase). As Schweidel et al. (2022) discuss, at each of these stages, consumers generate data through their interactions with other consumers and with brands. These digital signals provide the foundation for marketers to leverage AI-based solutions.

Drawing on the customer equity framework, we consider three distinct pillars that ultimately contribute to the value that customers provide to the firm: (1) customer acquisition, (2) relationship development, and (3) customer retention. At each of these stages, firms have access to different sources of data. At the customer acquisition stage, the information about a customer that the firm has available is more limited, as the firm must rely upon third-party information that has been tied to a particular prospect or an individual's prepurchase activities such as online browsing activity. Once customers have been acquired, firms can tie transactional activity, along with the response to marketing efforts, back to the same individual. This provides first-party data which can be used by the firm to personalize their marketing efforts to a greater extent. Such customization may be used to strengthen the relationship, which may contribute to cross-selling (e.g., Li, Sun, & Wilcox, 2005; Schweidel, Bradlow, & Fader, 2011), and to increased customer retention.

Extensive research in marketing has been conducted to identify *which individuals* should be targeted with marketing efforts. Rossi, McCulloch, and Allenby (1996) demonstrate how purchase histories could be used to tailor the promotions received by different consumers. Manchanda, Rossi, and Chintagunta (2004) show that physicians are not detailed optimally by pharmaceutical sales forces, with high-volume prescribers receiving more detailing visits than low volume prescribers, without regard for their sensitivity to detailing visits. Research has posited the use of customer lifetime value and customers' residual value (e.g., Dew & Ansari, 2018; Fader, Hardie, & Lee, 2005; Schmittlein, Morrison, & Colombo, 1987; Schweidel & Knox, 2013) as a means of scoring customers. Paralleling the findings of Manchanda et al. (2004), others have advocated for targeting decisions to be tied to the incremental impact of marketing efforts. Braun and Schweidel (2011) and Braun, Schweidel, and Stein (2015) suggest that the return on marketing, taking a customer's observed activity to date into account, should serve as the basis for targeting. Ascarza (2018) showed that firms should focus their efforts on retaining customers who were most sensitive to marketing interventions, regardless of their likelihood of churning. Lemmens and Gupta (2020) extend this idea further by incorporating incremental customer profitability due to the retention intervention.

Beyond the question of whom to target, marketers also face the decision of *when* they should send marketing communications to their customers. Using hidden Markov models, research has demonstrated that the timing of customer actions can be informative of future activity (e.g., Netzer, Lattin, & Srinivasan,

2008; Platzer & Reutterer, 2016; Schwartz, Bradlow, & Fader, 2014; Schweidel & Fader, 2009; Schweidel, Park, & Jamal, 2014; Zhang, Bradlow, & Small, 2015) Additionally, optimization methods have been used to determine the time at which marketing actions should be taken (e.g., Gönül & Shi, 1998; Kanuri, Chen, & Sridhar, 2018; Lewis, 2005; Li, Sun, & Montgomery, 2011; Montoya, Netzer, & Jedidi, 2010; Zhang, Kumar, & Cosguner, 2017).

While the questions of whom to target and when to target them have received ample attention in the marketing literature, research into the content of the marketing communications has been more limited. Bonfrer and Drèze (2009) developed a method for evaluating email campaign performance in real time, which can enable marketers to make changes to the creative content of emails that are underperforming. Hauser, Urban, Liberali, and Braun (2009) tailor the look and feel of a website to a visitor's cognitive style, inferred from their clickstream data. Schwartz, Bradlow, and Fader (2017) optimize the display ads shown on different websites over time, viewing the decision of which creative to show as a multiarm bandit problem. Kanuri et al. (2018) identify the optimal schedule for news content of different types to be posted.

To the best of our knowledge, targeting research focusing on whom to target, when, and with which content to employ has not yet been viewed through the lens of AI. However, there seems to be great potential. The resulting metrics and decision support systems, informed by prior data, could be incorporated into marketing automation, with the algorithms automating marketing decision-making in accordance with prespecified rules. In this light, firms are relying on AI solutions to mimic the behavior that would previously have been performed by marketers – selecting the appropriate individuals to receive marketing content, deciding on the creative elements of that content, and choosing when to distribute it.

Though empirical research has examined the way in which consumers respond to different types of content and used it to choose the optimal content to be used at the time of the next interaction with the customer, the extent to which research has informed content generation has been scant. Schwartz et al. (2017), for example, consider a finite number of display ads. Similarly, Kanuri et al. (2018) consider different types of content for a news provider. In both cases, content is essentially reduced to a multinomial choice problem.

There has been ample research investigating the elements of more successful content. Focusing on text-based communications, Berger and Milkman (2012) reported that consumers are more likely to share content that is highly arousing. Similarly, Rocklage, Rucker, and Nordgren (2021) find that emotional language in online reviews was diagnostic for market success of products. Berger and Packard (2018) found that cultural items that are atypical of their genre tend to be more popular, while Packard and Berger (2020) find that the use of second-person pronouns in cultural items results in more purchases. In addition to textual data, there has been research into the impact of images on market-relevant outcomes. Li and Xie (2020) showed that social media posts with images are more likely to be shared. Zhang, Lee, Singh, and Srinivasan (2021)

examined the relationship between image attributes and Airbnb property demand, identifying actionable insights for Airbnb hosts and photographers.

While research to date has identified characteristics associated with successful marketing content, most research has focused on a predefined set of hand-crafted features. Of this feature set, the ideal levels or combinations of levels can be identified. Consequently, extant marketing research can inform the development of new marketing content, whether it is textual copy or social media posts that consist of both text and images. However, since the feature set investigated in past research is not exhaustive, research does not offer insight into these missing components.

## 2. THE POTENTIAL FOR CONTENT GENERATION

As an illustration, let's consider a social media post. [Li and Xie's \(2020\)](#) research tells us that an image should be included to maximize consumer engagement. Drawing on [Zhang et al. \(2021\)](#), there are specific visual attributes that should have optimal levels. Beyond this guidance, what should be in the picture if the goal is to maximize engagement? We face a similar quandary with regard to the text of the post. While prior research suggests that we want to use arousing and emotional language (e.g., [Berger & Milkman, 2012](#); [Rocklage et al., 2021](#)) and second-person pronouns ([Packard & Berger, 2020](#)), what else should appear in the text?

Automated content generation has the potential to bridge the gap that exists between research that has identified specific content features that perform well and the creation and prediction of the performance of new marketing content. Recent advances in AI have yielded generative models which have found a few first pioneering applications in marketing. We next briefly review these models with potential for marketing content generation.

### *2.1 Generating Textual Content With Language Models*

With the increasing digitalization of our economy, and the accessibility of large amounts of digitized text (e.g., [Berger et al., 2020](#)), recent years have seen substantial advances in textual content generation using various AI architectures. These include long short-term memory (LSTM), convolutional, recurrent, and recursive neural networks ([Marchenko, Radyvonenko, Ignatova, Titarchuk, & Zhelezniakov, 2020](#)), and more recently also large-scale pretrained transformer systems (e.g., [Radford et al., 2019](#)). The basic intuition behind these methods is to predict the next likely unit (e.g., a word, a letter, or a token) by sampling from their known collection of units (their vocabulary), conditional on a given input sequence (e.g., of words) and on their model parameters, to fit the semantic context of the input sequence well.

Large-scale pretrained language models like GPT have proven to be superior over previous methods like LSTM on that task because of their innovative model architecture. This architecture essentially integrates word meaning and word

position information, a novel and efficient attention mechanism (i.e., to put varying degrees of attention on a given set of word inputs; Vaswani et al., 2017), and feed forward neural networks over several transformer decoder layers to recreate natural language patterns.<sup>4</sup> Unlike previous NLG models, transformers are capable of producing convincing, human-looking, and coherent textual passages (Floridi & Chiriatti, 2020), while the computational part is no more complex than a massive amount of matrix multiplications. Taken together, this yields well-performing open-source transformer models that are broadly accessible for both research and commercial purposes and use-cases, pretrained on large amounts of unsupervised textual data.

While pretrained models are broadly applicable and produce content with a common and general linguistic structure not tailored for a specific context, many marketers might have specific goals for content generation in mind. For example, marketing managers might aim at generating content for social media postings, coming with the necessity of very specific optimal textual features like specific topics, styles, or patterns like the usage of arousing language and of second-person pronouns. Technically, there are three broad ways (or a combination of these) to integrate such features in content generation:

First, one might adapt the transformer's model parameters, either via a complete retraining (i.e., a fresh start) or in a process called fine-tuning. In fine-tuning, we use the pretrained model parameters as a starting point to build on the transformer's general language knowledge, to further train the model on domain specific content like websites, or social media postings for integrating domain specific optimal features. For example, fine-tuning a pretrained transformer on website content enables the transformer to generate new website content being reflective of the topics, styles, etc., commonly appearing on websites (e.g., Reisenbichler, Reutterer, Schweidel, & Daniel, 2022), fine-tuning on social media postings enables transformers to generate content in the style of social media posts, and so on.

Second, one can adapt the transformer's model architecture in order to adjust the prevalence of certain text attributes or sentiment. A recent example among several of such approaches is plug and play language models (PPLM; Dathathri et al., 2019), an attribute-enriched GPT-2 model that in essence alters GPT-2's output prediction of the next unit by increasing the probability for words that appear in customized word lists to feed in. This allows for directing content generation on a fine-grained level, for example, by feeding in a word list containing words that are associated with arousal, and a word list containing second-person pronouns for generating social media posts that feature these attributes.

Third, one might construct a content selection method that measures optimal feature integration (e.g., Reisenbichler et al., 2022). Due to the probabilistic content generation process, generating several social media content samples could yield content that contains desirable features such as language consistent with the brand tone, or undesirable features such as racist language. Thus, one might opt for automatically selecting the most optimal content samples based on a measure of the integration of optimal features in the generated content.

Which route for feature integration to follow depends on the specific problem, resources, and capabilities of researchers and businesses. However, it is noteworthy that combining two or several of the above-mentioned three ways for feature integration might be beneficial for specific applications (e.g., [Reisenbichler et al., 2022](#)).

## 2.2 Generating Synthetic Images

Firms use various types of information to show their products and services, engage their existing customers, and attract new users. Among these, images are an efficient channel to convey information to consumers, resulting in firms undertaking effort to design and manipulate images that will yield their desired outcomes. While undertaking such a task manually is costly and not scalable, recent advances in machine learning offer new opportunities for marketers. In particular, the unsupervised (i.e., learning the regularities or patterns from data in such a way that human intervention is not needed) generative models have been extensively explored, developed, and validated.

*Approaches for generating visual content.* One approach to image generation builds off the previously discussed language models. Language models create a vector representation of the text. Images can be also projected into an embedding space as vectors. Through the training process where a set of images and texts are paired, we can map two embedding spaces, which can consequently lead to the generation of photorealistic images, such as “an illustration of a baby hedgehog in a Christmas sweater walking a dog” or “the exact same cat on the top as a sketch on the bottom” ([Ramesh et al., 2021](#)).

Generative adversarial networks (GANs) represent a methodological approach that, among others, has also been proposed as a means by which visual content can be generated. First proposed by [Goodfellow et al. \(2014\)](#), GANs frame the problem as a supervised learning task with two sub-networks: the generator network (G) that we train to produce new examples (e.g., images) and the discriminator network (D) that aims to classify examples as either real (from the domain) or generated. The two sub-networks (two players) are trained jointly in a min-max game, where G aims to fool the D by producing close-to-realistic examples and D becomes better in distinguishing between real and generated data. The game continues until the discriminator network D is fooled about half the time, indicating that the generator network G is generating plausible examples.

After the inception of GANs, variants have been explored and developed to tackle specific problems. [Mirza and Osindero \(2014\)](#) proposed a conditional GAN (CGAN) to incorporate information such as product categories and image labels into the two sub-networks. [Karras, Laine, and Aila \(2018\)](#) proposed a style-based GAN (StyleGAN) where the generator network controls the image synthesis through scale-specific amendments to the styles without compromising the generation quality. To do so, StyleGAN divides input features into three types: (i) coarse features (e.g., face shape), (ii) medium features (e.g., finer facial

features such as the hair style and eyes open/closed), and (iii) fine-grained features (e.g., high-resolution microfeatures like the hair color).

Leveraging more data into large GAN models to capture highly complex patterns for better image generation has attracted lots of attention. However, it also carries an intense computational burden. Fortunately, given the availability of large-scale computing resources (e.g., graphical processing unit or tensor processing unit), researchers have begun to develop high-quality generative models by training large GAN models that make use of many neural network parameters. Big GAN (BigGAN), one of the best present representatives of GAN models, was proposed by [Brock, Donahue, and Simonyan \(2018\)](#) and has shown outstanding performance in large and high-fidelity image generation, with the limitation that the diversity of generated images is much lower than that of real images with the same size.

*Applications of visual content generation:* [Goodfellow et al. \(2014\)](#) demonstrated the initial use of GANs to generate examples of handwritten digits and human faces. However, GANs also have the potential to generate more complex images. For example, [Radford, Metz, and Chintala \(2015\)](#) developed a deep convolutional GAN and demonstrated its ability to generate images of bedrooms. [Karras, Aila, Laine, and Lehtinen \(2017\)](#) showed the ability of PGGAN to produce real-looking celebrity faces, as well as objects and scenes. BigGAN was used to generate images that are practically indistinguishable from real ones.

The second typical application of GANs for image generation are translation tasks. [Isola, Zhu, Zhou, and Efros \(2017\)](#) developed a pix2pix approach for many image-to-image translation tasks, such as semantic images to photographs of cityscapes and buildings, satellite photographs to Google maps, photos from day to night, black and white images to colorful ones, and sketches to color photos. CycleGAN by [Zhu, Park, Isola, and Efros \(2017\)](#) illustrated use cases including the translation of images to artistic paintings, horses to zebras, and photographs from summer to winter. Such image-to-image translation can also be found using CGAN ([Isola et al., 2017](#)).

In addition to translation tasks, GANs have also been used to edit images. For example, [Perarnau, Van De Weijer, Raducanu, and Álvarez \(2016\)](#) developed their own model, IcGAN, to reconstruct photos of faces with specific features, such as changes in hair color, style, facial expression, and even gender. [Brock, Lim, Ritchie, and Weston \(2016\)](#) presented a face photo editor using a hybrid of a variational autoencoder and GANs to modify human faces with certain features. GANs have also been used to age faces, generating photographs of an individual at different apparent ages, from younger to older ([Antipov, Baccouche, & Dugelay, 2017](#); [Zhang, Xu et al., 2017](#)). Photo blending has also become practical, allowing for the incorporation of elements from different images such as fields, mountains, and other large structures ([Wu et al., 2017](#)).

*Image generation applications for marketing.* Turning our attention to promising potential applications within marketing, one direction that can be pursued is the personalization of images that are used in marketing communications such as emails and social media posts. [Liu, Dyzabura, and Mizik \(2019\)](#) demonstrate how deep learning can be used to examine brand personalities as expressed social



media posts from consumers which have tagged specific brands. Taking this work in a slightly different direction, consumers' expressed personality could be inferred from the set of images that they have posted to social media platforms. Rather than developing entirely new visual content for each consumer, brands may seek to leverage image generation to modify visuals, tweaking them so that they align more with a consumer's expressed personality.

Another direction that may be of interest to pursue is generating images that will complement existing text content. Doing so would build upon prior morphing research on website design (e.g., [Hauser et al., 2009](#)). Rather than generating an image that is expected to perform well on its own, by taking into account the textual content with which it will be published could result in the production of content that yields higher conversion rates. Such approaches may also incorporate consumers' individual preferences.

We can also envision the use of image generation to support social media campaigns. With brands frequently posting content across multiple platforms, AI-based solutions could provide guidance as to the content of different posts, composed of text and images, to be posted across different platforms and when such content should be posted.

One challenge that brands may face is the computational efficiency associated with content generation. Rather than attempting to conduct tasks such as personalization on the fly, firms will likely face technical limitations. While customization based on user profiles to support digital marketing efforts such as email campaigns need not be conducted in real time, other efforts such as personalized content on landing pages will require identifying the extent to which intensive calculations can be performed offline. Doing so will alleviate marketers of the need to perform such calculations on a real-time basis.

### **3. SUPPORTING CUSTOMER EQUITY MANAGEMENT WITH CONTENT GENERATION**

Before examining the ways in which AI-driven content generation can support customer equity, we first describe a general approach by which generative AI can be used. To ensure that AI-generated content is desirable, it is necessary to identify content features that have been linked to success and develop a scoring model. In the case of search engine optimization (SEO), this may be based on whether or not content ranks on the first page of search results (e.g., [Reisenbichler et al., 2022](#)). For social media posts, a scoring model may be based on the volume of engagement or the number of shares a post receives ([Berger & Milkman, 2012](#); [Tellis, MacInnis, Tirunillai, & Zhang, 2019](#)).

Armed with a scoring model against which the AI-generated content may be evaluated, marketers may then use generative models to develop candidate content. This might rely entirely on pretrained models or may involve the fine-tuning of pretrained models, so that they are more directly applicable to a specific marketing domain. The candidate content can then be evaluated using the scoring model. Once the top performing candidates have been identified, the

AI-generated content can be provided to a human decision-maker for potential revision. At a minimum, this final step is needed to ensure factual accuracy but can also be used to ensure that content is consistent with a brand's guidelines.

We next turn our attention to each of the components of customer equity and identify specific opportunities in which marketers can employ content generation.

### *3.1 Customer Acquisition*

*Online search:* One of the means by which marketers may attract new customers is through SEO. For consumers who turn to search engines, SEO is concerned with garnering a high ranking in the search results. Consumers are more likely to consider the search results that appear earlier in the organic listings. Rankings that appear higher in the organic results are of higher quality (Berman & Katona, 2013), based on factors including the content (e.g., Liu & Toubia, 2018) and the link structure.

One way in which AI can support firms' SEO efforts is by identifying the appropriate keywords for which to develop content. Prior to developing search engine-optimized content, firms undertake keyword research to ensure that they develop content for the search terms used by consumers. Jerath, Ma, and Park (2014) probed the keywords used to conduct the search and found that consumers using less popular keywords were more likely to click on search results and are closer to making a purchase. While consumers use the less popular keywords less frequently, they are more likely to make purchases following queries that use these terms. Therefore, it is important to not only have content that performs well for popular keywords but also to identify such long tail keywords that are strongly associated with purchasing. Website traffic data can support analyses akin to that conducted by Jerath et al. (2014) to identify the set of keywords that are important for a given firm. While an examination of website traffic allows for the identification of the keywords bringing traffic to a given website, AI can also be used to examine keywords of relevance to other websites based on their search engine performance. In doing so, firms may expand the list of keywords for which they develop content beyond the current set.

Once the keywords have been identified, AI can be used to generate search engine-optimized content. Reisenbichler et al. (2022) demonstrate how pre-trained language models can be fine-tuned to produce such content. In particular, the authors fine-tune a pretrained language model with the content of the web pages that appear on the first page of Google search results. That is, recognizing that the top 10 search results as having the desirable content according to Google's search engine algorithm, they update the pretrained language model to generate content based on the linguistic patterns of these 10 pages. Comparing the generated content, after it is edited by a human, to the content that was generated by SEO specialists, they demonstrate that AI coupled with a "human-in-the-loop" can not only reduce content costs significantly but also improve performance.

In addition to generating content designed to attract visitors through organic search results, marketers have also focused on what they may do to perform better in paid search. Yao and Mela (2011) showed that those who frequently click on sponsored search results put an emphasis on the position of the advertisement, and that approximately 10% of consumers are responsible for 90% of clicks. They further demonstrate the value associated with being able to vary bids based on consumer segment, demonstrating increases in both advertiser revenue and consumer welfare. Exploring the use of branded and nonbranded keywords, Blake, Nosko, and Tadelis (2015) find that branded keywords do not have a measurable short-term impact. For nonbranded keywords, however, infrequent users respond positively to paid search, while frequent users do not, resulting in net negative average returns.

Rutz, Bucklin, and Sonnier (2012) examine keyword conversion in paid search, recognizing that the position of the advertisement is endogenous. They find that higher rankings among the paid search results are associated with both increased click-through rates and higher conversion rates. Importantly, they demonstrate that producing a keyword list using their estimates outperforms the approach that relies on observed conversion and click-throughs.

One of the ways in which AI can support customer acquisition via paid search is by generating textual advertising content. There have been efforts to predict the performance of paid search advertisements based on the content of the advertisements (e.g., Rutz, Sonnier, & Trusov, 2017). Bhattacharya, Gong, and Wattal (2021), for example, examine the performance of ad content in a competitive context. They find that when attempting to bid on a high-quality brand's keywords, ad content that highlights vertical differentiation is more effective. In contrast, when attempting to bid on the keywords of low-quality brands, messages featuring horizontal differentiation outperform control ads. Though the nature of differentiation relative to the competition is one aspect of the text, there are additional textual factors. Moreover, the content of the text and bid amounts may interact. The use of AI employing generative models to develop ad copy, coupled with carefully designed experiments, could support acquisition efforts through paid search.

*Social media messaging:* Much research has probed those factors that are related to the virality and persuasiveness of social media messages. Among the messaging factors that have been considered are emotions (e.g., Akpınar & Berger, 2017; MacInnis, Rao, & Weiss, 2002; Tellis et al., 2019), arousal (e.g., Berger & Milkman, 2012), explicit versus implicit endorsements (e.g., Packard & Berger, 2017), and trust cues (e.g., Packard, Gershoff, & Wooten, 2016).

In addition to the textual information, the visuals used in social media messages may also play a role, such as the prominence of the brand (e.g., Pieters & Wedel, 2004) and visual complexity (e.g., Pieters, Wedel, & Batra, 2010). Moving beyond specific features, Liu, Dzyabura, and Mizik (2020) developed a deep learning model to identify brand perceptions associated with images posted on social media. Similarly, Hartmann, Heitmann, Schamp, and Netzer (2021) employ deep learning to examine the performance of different types of visual posts on social media, finding that consumer selfies that feature brands receive

more likes and comments, but that brands selfies (that do not feature consumers) garner higher levels of purchase intentions.

AI tools could be used in developing social media messages that are engineered to perform well on social media (e.g., Kumar, Bhaskaran, Mirchandani, & Shah, 2013; Wang & Kim, 2017). The text and images generated to produce the social media messages could be tailored to different consumer segments, taking advantage of user profiles that may be developed based on their social media activity (e.g., Adamopoulos, Ghose, & Todri, 2018; Zhang, Moe, & Schweidel, 2017). These posts may prominently feature the brand's logo, which may evoke certain brand perceptions and could be generated using AI (e.g., Dew, Ansari, & Toubia, 2021). Depending on the intended recipient of the message, different visuals may be produced to evoke specific brand perceptions (e.g., Liu et al., 2020).

### *3.2 Relationship Development*

Once customers have been acquired, the role of marketing content shifts. Rather than building awareness, the intent becomes to strengthen the relationship and remain top-of-mind for customers vis-à-vis competitors. Doing so requires that firms understand who their customers are, to be able to identify their needs from the data available, and to design communications that address these needs accordingly. A number of recent methodological advances that enables marketers to “mine” customer–firm interactions and online conversations (e.g., Lee & Bradlow, 2011; Liu, Lee, & Srinivasan, 2019; Netzer, Feldman, Goldenberg, & Fresko, 2012) can help for this purpose. For example, Timoshenko and Hauser (2019) propose a deep convolutional network approach that enables firms to identify customer needs from online reviews. Gabel, Guhl, and Klapper (2019) use a neural network–based method to map market structures based on market basket data, while Yang, Zhang, and Kannan (2021) propose a deep network representation learning framework to make inferences on competitive market structures based on social media brand engagement data. Such AI-supported tools not only help to condense and summarize complex interdependencies among brands and concepts but also have the potential to support marketers in developing and communicating new creative content to their target groups.

The value of historic data to inform personalized marketing efforts has been well established, from couponing (e.g., Gabel & Timoshenko, 2021; Rossi et al., 1996) to website design (e.g., Hauser et al., 2009; Urban, Liberali, MacDonald, Bordley, & Hauser, 2014). Advances in text and image analytics have paved the way for marketers to learn about their customers based on the content that they have previously created on social media. For example, based on user activity on Twitter, Zhang, Moe, and Schweidel (2017) profile individuals based on the topics of their recent activity, yielding an interest profile for each user. The authors further demonstrate that individuals are more likely to share content that is topically consistent with their social media–derived interest profile. Once such individual profiles have been derived, whether it be based on language, social

connections (e.g., [Culotta & Cutler, 2016](#)), or brand preferences, the content that is generated can be customized based on the individual's profile.

Deploying such personalized profiles may take on multiple forms. The content of landing pages could be customized based on the individual profiles. Similarly, social media messaging (both organic and sponsored) can be tailored to the individual. Firms can also tailor the content of email messages based on what has been found to perform well historically with the customer. In a manner similar to the approach taken by [Reisenbichler et al. \(2022\)](#) for generating website content for SEO purposes, text could be produced. But, rather than the same text being used for the marketing to all individuals, individual information would be incorporated to allow for personalization based on the historic information available about the customer and the relationship between the customer and firm (e.g., [Schweidel et al., 2022](#)).

However, textual content is just one part of the puzzle. Visual content plays an important role in engaging customers. Research has identified visual elements that are associated with increased virality and engagement (e.g., [Li & Xie, 2020](#); [Tellis et al., 2019](#); [Zhang et al., 2021](#)). [Liu et al. \(2020\)](#) demonstrated how a deep learning model could be used to identify how a social media image associated with a brand would be perceived. Such models may facilitate the development of generative models for images, such as DALL-E and GauGAN, that allow for the production of visuals that are similar to the visual style to which customers have positively responded.<sup>5</sup> In addition to the generation of individualized images, by developing visual profiles of customer preferences, AI can support the selection of the optimal image from an existing set for each customer.

Video content can also be generated, taking into account customer-level information. [Liu, Shi, Teixeira, and Wedel \(2018\)](#) demonstrated how real-time emotional reactions to trailers could be used to inform which clips from the movie could be used to market the film. Though the authors' proposed approach is to develop an optimal trailer for a given film, the method they develop could be generalized to develop content tailored to customer segments, or perhaps individual customers. For such approaches to succeed, each individual's (or segment's) content preferences would be learned through responses to prior content to which they had been exposed. [Rajaram and Manchanda \(2020\)](#) investigate the effect of the nature of advertising content in influencer videos on viewing behavior, interaction rates, and sentiment. Using a deep learning architecture with transfer learning for different modalities (text, audio, and images), they find that brand name inclusion, human sounds, and music in the early phases of video content are associated with beneficial managerial outcomes.

Though the use of AI to support relationship development is tempting, research advises that care must be taken in its deployment. [Longoni and Cian \(2020\)](#) demonstrate that consumers react more favorably to AI-based recommendations when it comes to utilitarian attributes, but that human recommendations are preferred for hedonic attributes. [Luo, Tong, Fang, and Qu \(2019\)](#) similarly documented adverse reactions to AI, this time in the case of chatbots compared to live agents. The authors find that purchase rates decrease by nearly 80% when the identity of a chatbot is disclosed compared with when its identity is not

disclosed. Similarly, [Crolic, Thomaz, Hadi, and Stephen \(2022\)](#) demonstrate in a series of studies that human-like chatbot conversations have a negative effect on customer satisfaction and purchase intention if customers are in an angry state (e.g., because of a negative customer experience with the firm). Taken together, these findings suggest that there may be limits on the extent to which content that has been produced by AI can be leveraged, with a “human-in-the-loop” perhaps being necessary to ensure the appropriateness of the content (e.g., [Reisenbichler et al., 2022](#)).

### *3.3 Customer Retention*

Just as AI can be used to support customer acquisition and relationship development, it has the potential to improve efforts to increase customer retention. One way in which AI tools can support customer retention is through the identification of individuals who would and would not benefit from being targeted with retention efforts. [Braun and Schweidel \(2011\)](#) and [Lemmens and Gupta \(2020\)](#) advocate for selecting targets for retention efforts based on the financial impact of retention interventions. [Godinho de Matos, Ferreira, and Belo \(2018\)](#), [Nitzan and Libai \(2011\)](#) and [Ascarza, Ebbes, Netzer, and Danielson \(2017\)](#) demonstrate that social effects should be taken into account in managing the relationships, with social ties affecting an individual’s tendency to churn.

Beyond identifying who should be targeted, research has also addressed the question of when individuals should be targeted with communications. [Schweidel et al. \(2011\)](#) show that the tendency to churn evolves over the course of the relationship and may exhibit a nonmonotonic pattern. [Ma, Sun, and Kekre \(2015\)](#) demonstrate that customers’ social media activity is indicative of their tendency to churn. [Ascarza, Netzer, and Hardie \(2018\)](#) distinguish between overt and silent churners. They find that higher levels of customer engagement, such as responding to marketing communications, are associated with higher levels of observed churn, while silent churn tends to occur early in the relationship. As such research has revealed, dynamic models update customers’ churn tendencies based on the behavior observed to date, supporting the targeting of the right individual at the right time with retention efforts. In a direct marketing context, [Sarkar and De Bruyn \(2021\)](#) showed that recurrent neural networks can automatically identify predictive features directly from the flow of raw data. More recently, [Valentin, Reutterer, Platzer, and Kalcher \(2022\)](#) proposed a self-supervised sequence-to-sequence model that leverages the automatic feature detection capabilities of such neural network architectures to make long-term customer predictions and to accurately detect customers at risk.

However, once the right customers have been identified for retention efforts, what content should retention efforts employ? As with customer acquisition and relationship development, the heterogeneity that exists in the customer bases all, but guarantees that there will not be a one-size-fits-all solution. With the exception of [Ma et al. \(2015\)](#), there has been scant empirical research that has considered the content used by firms in an effort to retain customers. One dimension that should be considered in developing content for retention

interventions is the medium that is used (e.g., email vs. social media messages). In addition to the medium, the content of the message, which can be informed by the content with which customers have previously interacted, from emails and social media posts to website browsing, content preferences can be identified. Taken together, candidate retention interventions can be developed. The efficacy of these efforts, however, will depend on whom they are targeted and when they are deployed. Retention efforts should simultaneously consider all four of these decisions – whom to communicate with, when to make the effort, how to reach them, and what to say. While preliminary efforts have approached these one at a time or in pairs, AI may offer a means of generating and testing candidate retention messages targeted at different subsets of the customer base. Incorporated into a dynamic framework, the learnings from the AI-supported deployment of retention efforts can subsequently inform later efforts.

## 4. CONSIDERATIONS FOR THE USE OF AI-SUPPORTED CONTENT GENERATION

### *4.1 Consumer Reactions*

Technologies' inevitable march forward will result in the advancement of generative models. Several businesses have been developed leveraging large language models to streamline marketing tasks, such as the production of social media and website text. Generative video models have been used to create advertising content for local markets at a fraction of the cost when compared to producing the advertisement from scratch.<sup>6</sup> And while marketing use cases abound, how will consumers react to such content?

As Luo et al. (2019) and Longoni and Cian (2020) revealed, if consumers are aware of the fact that they are interacting with AI-generated content, the reaction may not be favorable. Such preferences may evolve over time as consumers become more accustomed to interacting with AI through everyday technologies such as autocompletion for email and text messaging, and digital voice assistants. Part of consumers' reluctance to accept AI-generated content may also have to do with the specific context. While consumers may be comfortable interacting with an AI customer service agent to make a reservation, they may prefer to engage with a human for the reporting of sensitive information such as medical results. Despite consumer hesitance, performance may be enhanced through the use of AI (e.g., Longoni, Bonezzi, & Morewedge, 2019), making it critical that consumers understand the technology with which they are interacting (Cadario, Longoni, & Morewedge, 2021). Having trust in the organizations deploying it may also affect consumers' eventual acceptance of the technology.

### *4.2 Potential Abuse and the Need for Regulation*

Beyond consumers' acceptance of the technology, we must also ask ourselves about the potential for the technology to be abused. Already, we have seen the use of generative video models to develop "deepfakes," which could be used to

defame public figures by placing them in inappropriate contexts or altering their statements.<sup>7</sup> AI-supported content generation enables the dissemination of misinformation at scale, whether it be brand-related misinformation or misinformation related to public issues such as a global health emergency or elections. There is an urgent need not only for the regulation on the use of such technology by governments; given the pace of innovation, however, it is equally if not more urgent that digital platforms be involved in curbing the inappropriate use of AI-generated content. Doing so will require the ability to detect content that has been generated by AI, ideally in an automated fashion, so that such countermeasures may be deployed at scale.

In addition to needing to ensure the appropriate use of AI-supported content generation, there are also concerns as to its perpetuation of existing biases. Large language models, and generative models more broadly, have been referred to as “stochastic parrots” (Bender, Gebru, McMillan-Major, & Shmitchell, 2021). While the content generated by these models may give the appearance of human understanding, they are based on the data on which they are trained, yielding content that follows the same patterns. A recognized consequence of this is the reinforcement of existing biases, such as those related to gender and racial differences. For marketers that intend to use these generative models, care must be taken to ensure that such biases are not unintentionally propagated. This may involve evaluating AI-generated content, and marketing content more broadly, for signs of bias.

### *4.3 Workforce Implications*

While the dangers to society stemming from the misinformation and biases perpetuated by AI are apparent, a further consideration of which we must be mindful are the threats posed by the generation of AI-supported content to the workforce. The promise of increased efficiency is contributing to the adoption of AI in business (see, Reisenbichler et al., 2022). AI tools can be used for a range of content-related marketing tasks, such as automating the posting of social media messages that have been crafted by human contributors or using AI to create the first draft of creative content. In doing so, less time and consequently fewer employees are needed for content-related marketing functions. While AI may take over more analytical (and thus, automatable) tasks in such scenarios, the “softer” and more intuitive or empathic tasks will remain the domain for employees (cf. Huang & Rust, 2018).

Though automation through the adoption of AI poses a threat to the existing workforce, there will also be opportunities for employment created by AI. For example, the generative models that we adopt must be developed and evaluated. In addition to specifying the evaluation criteria, such as quantifying the extent of bias, there will be a need to test the underlying models. If we were to view AI-supported content generation as a creative aid rather than a replacement, the adoption of AI may make today’s marketers more effective, leading to business growth that creates new employment opportunities. In evaluating the totality of



AI-supported content generation's impact on workers, it is important to bear in mind both the positive and negative effects that such technologies may have.

## 5. CONCLUSION

Rooted in the customer equity framework, we discuss the existing research that have paved the way for the use of AI-supported content generation in marketing. Customer acquisition, relationship development, and customer retention efforts can be deployed that make use of content that has been generated based on what has previously been successful for a given endeavor. Taking this a step further, AI-generated content can be tailored to the segment or individual customers, ushering in true one-to-one marketing on a scale not previously experienced.

## NOTES

1. <https://www.microsoft.com/en-us/research/blog/using-deepspeed-and-megatron-to-train-megatron-turing-nlg-530b-the-worlds-largest-and-most-powerful-generative-language-model/>
2. <https://openai.com/dall-e-2/>
3. <https://www.nvidia.com/>
4. For a more detailed technical description of the transformer system GPT-2, we refer the reader to [Reisenbichler et al. \(2022\)](#).
5. <https://www.zdnet.com/article/the-absurd-beauty-of-hacking-nvidias-gaugan-2-ai-image-machine/>
6. <https://www.cbsnews.com/news/deepfake-artificial-intelligence-60-minutes-2021-10-10/>
7. <https://www.reuters.com/technology/deepfake-anyone-ai-synthetic-media-tech-enters-perilous-phase-2021-12-13/>

## REFERENCES

- Adamopoulos, P., Ghose, A., & Todri, V. (2018). The impact of user personality traits on word of mouth: Text-mining social media platforms. *Information Systems Research*, 29(3), 612–640.
- Akpınar, E., & Berger, J. (2017). Valuable virality. *Journal of Marketing Research*, 54(2), 318–330.
- Antipov, G., Baccouche, M., & Dugelay, J. L. (2017, September). Face aging with conditional generative adversarial networks. In *2017 IEEE international conference on image processing (ICIP)* (pp. 2089–2093). IEEE.
- Ascarza, E. (2018). Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research*, 55(1), 80–98.
- Ascarza, E., Ebbes, P., Netzer, O., & Danielson, M. (2017). Beyond the target customer: Social effects of customer relationship management campaigns. *Journal of Marketing Research*, 54(3), 347–363.
- Ascarza, E., Netzer, O., & Hardie, B. G. (2018). Some customers would rather leave without saying goodbye. *Marketing Science*, 37(1), 54–77.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency* (pp. 610–623).
- Berger, J., Humphreys, A., Ludwig, S., Moe, W. W., Netzer, O., & Schweidel, D. A. (2020). Uniting the tribes: Using text for marketing insight. *Journal of Marketing*, 84(1), 1–25.
- Berger, J., & Milkman, K. L. (2012). What makes online content viral? *Journal of Marketing Research*, 49(2), 192–205.

- Berger, J., & Packard, G. (2018). Are atypical things more popular? *Psychological Science*, 29(7), 1178–1184.
- Berman, R., & Katona, Z. (2013). The role of search engine optimization in search marketing. *Marketing Science*, 32(4), 644–651.
- Bhattacharya, S., Gong, J., & Wattal, S. (2021). Competitive poaching in search advertising: Two randomized field experiments. *Information Systems Research*.
- Blake, T., Nosko, C., & Tadelis, S. (2015). Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica*, 83(1), 155–174.
- Bonfrer, A., & Drèze, X. (2009). Real-time evaluation of e-mail campaign performance. *Marketing Science*, 28(2), 251–263.
- Braun, M., & Schweidel, D. A. (2011). Modeling customer lifetimes with multiple causes of churn. *Marketing Science*, 30(5), 881–902.
- Braun, M., Schweidel, D. A., & Stein, E. (2015). Transaction attributes and customer valuation. *Journal of Marketing Research*, 52(6), 848–864.
- Brock, A., Donahue, J., & Simonyan, K. (2018). Large scale GAN training for high fidelity natural image synthesis. arXiv preprint arXiv:1809.11096.
- Brock, A., Lim, T., Ritchie, J. M., & Weston, N. (2016). Neural photo editing with introspective adversarial networks. arXiv preprint arXiv:1609.07093.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... Amodei, D. (2020). Language models are few-shot learners. arXiv preprint arXiv:2005.14165.
- Cadario, R., Longoni, C., & Morewedge, C. K. (2021). Understanding, explaining, and utilizing medical artificial intelligence. *Nature Human Behaviour*. doi:10.1038/s41562-021-01146-0
- Crolic, C., Thomaz, F., Hadi, R., & Stephen, A. T. (2022). Blame the bot: Anthropomorphism and anger in customer–chatbot interactions. *Journal of Marketing*, 86(1), 132–148.
- Culotta, A., & Cutler, J. (2016). Mining brand perceptions from twitter social networks. *Marketing Science*, 35(3), 343–362.
- Dathathri, S., Madotto, A., Lan, J., Hung, J., Frank, E., Molino, P., ... Liu, R. (2019). Plug and play language models: A simple approach to controlled text generation. arXiv preprint arXiv:1912.02164.
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
- Dew, R., & Ansari, A. (2018). Bayesian nonparametric customer base analysis with model-based visualizations. *Marketing Science*, 37(2), 216–235.
- Dew, R., Ansari, A., & Toubia, O. (2021). Letting logos speak: Leveraging multiview representation learning for data-driven branding and logo design. *Marketing Science*, 41(2), 401–425.
- Fader, P. S., Hardie, B. G., & Lee, K. L. (2005). RFM and CLV: Using iso-value curves for customer base analysis. *Journal of Marketing Research*, 42(4), 415–430.
- Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30(4), 681–694.
- Gabel, S., Guhl, D., & Klapper, D. (2019). P2V-MAP: Mapping market structures for large retail assortments. *Journal of Marketing Research*, 56(4), 557–580.
- Gabel, S., & Timoshenko, A. (2021). Product choice with large assortments: A scalable deep-learning model. *Management Science*.
- Godinho de Matos, M., Ferreira, P., & Belo, R. (2018). Target the ego or target the group: Evidence from a randomized experiment in proactive churn management. *Marketing Science*, 37(5), 793–811.
- Gönül, F., & Shi, M. Z. (1998). Optimal mailing of catalogs: A new methodology using estimable structural dynamic programming models. *Management Science*, 44(9), 1249–1262.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27.
- Gupta, S., & Lehmann, D. (2005). *Managing customers as investments: The strategic value of customers in the long run*. Upper Saddle River, NJ: Wharton School Publishing.
- Hartmann, J., Heitmann, M., Schamp, C., & Netzer, O. (2021). The power of brand selfies. *Journal of Marketing Research*, 58(6), 1159–1177.

- Hauser, J. R., Urban, G. L., Liberali, G., & Braun, M. (2009). Website morphing. *Marketing Science*, 28(2), 202–223.
- Huang, M. H., & Rust, R. T. (2018). Artificial intelligence in service. *Journal of Service Research*, 21(2), 155–172.
- Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 5967–5976) doi:10.1109/CVPR.2017.632
- Jerath, K., Ma, L., & Park, Y. H. (2014). Consumer click behavior at a search engine: The role of keyword popularity. *Journal of Marketing Research*, 51(4), 480–486.
- Kanuri, V. K., Chen, Y., & Sridhar, S. (2018). Scheduling content on social media: Theory, evidence, and application. *Journal of Marketing*, 82(6), 89–108.
- Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2017). Progressive growing of GANs for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196.
- Karras, T., Laine, S., & Aila, T. (2018). A style-based generator architecture for generative adversarial networks. arXiv preprint arXiv:1812.04948.
- Kumar, V., Bhaskaran, V., Mirchandani, R., & Shah, M. (2013). Practice prize winner—Creating a measurable social media marketing strategy: Increasing the value and ROI of intangibles and tangibles for hokey pokey. *Marketing Science*, 32(2), 194–212.
- Lee, T. Y., & Bradlow, E. T. (2011). Automated marketing research using online customer reviews. *Journal of Marketing Research*, 48(5), 881–894.
- Lemmens, A., & Gupta, S. (2020). Managing churn to maximize profits. *Marketing Science*, 39(5), 956–973.
- Lemon, K. N., & Verhoef, P. C. (2016). Understanding customer experience throughout the customer journey. *Journal of Marketing*, 80(6), 69–96.
- Lewis, M. (2005). Incorporating strategic consumer behavior into customer valuation. *Journal of Marketing*, 69(4), 230–238.
- Li, S., Sun, B., & Montgomery, A. L. (2011). Cross-selling the right product to the right customer at the right time. *Journal of Marketing Research*, 48(4), 683–700.
- Li, S., Sun, B., & Wilcox, R. T. (2005). Cross-selling sequentially ordered products: An application to consumer banking services. *Journal of Marketing Research* 42(2), 233–239.
- Liu, L., Dzyabura, D., & Mizik, N. (2020). Visual listening in: Extracting brand image portrayed on social media. *Marketing Science*, 39(4), 669–686.
- Liu, X., Lee, D., & Srinivasan, K. (2019). Large-scale cross-category analysis of consumer review content on sales conversion leveraging deep learning. *Journal of Marketing Research*, 56(6), 918–943.
- Liu, X., Shi, S. W., Teixeira, T., & Wedel, M. (2018). Video content marketing: The making of clips. *Journal of Marketing*, 82(4), 86–101.
- Liu, J., & Toubia, O. (2018). A semantic approach for estimating consumer content preferences from online search queries. *Marketing Science*, 37(6), 930–952.
- 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), 1–19.
- Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research*, 46(4), 629–650.
- Longoni, C., & Cian, L. (2020). Artificial intelligence in utilitarian vs. hedonic contexts: The “word-of-machine” effect. *Journal of Marketing*. doi:10.1177/0022242920957347
- Luo, X., Tong, S., Fang, Z., & Qu, Z. (2019). Frontiers: Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases. *Marketing Science*, 38(6), 937–947.
- MacInnis, D. J., Rao, A. G., & Weiss, A. M. (2002). Assessing when increased media weight of real-world advertisements helps sales. *Journal of Marketing Research*, 39(4), 391–407.
- Manchanda, P., Rossi, P. E., & Chintagunta, P. K. (2004). Response modeling with nonrandom marketing-mix variables. *Journal of Marketing Research*, 41(4), 467–478.
- Marchenko, O. O., Radyvonenko, O. S., Ignatova, T. S., Titarchuk, P. V., & Zhelezniakov, D. V. (2020). Improving text generation through introducing coherence metrics. *Cybernetics and Systems Analysis*, 56(1), 13–21.

- Ma, L., Sun, B., & Kekre, S. (2015). The squeaky wheel gets the grease—An empirical analysis of customer voice and firm intervention on Twitter. *Marketing Science*, 34(5), 627–645.
- Mirza, M., & Osindero, S. (2014). Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784.
- Montoya, R., Netzer, O., & Jedidi, K. (2010). Dynamic allocation of pharmaceutical detailing and sampling for long-term profitability. *Marketing Science*, 29(5), 909–924.
- Netzer, O., Feldman, R., Goldenberg, J., & Fresko, M. (2012). Mine your own business: Market-structure surveillance through text mining. *Marketing Science*, 31(3), 521–543.
- Netzer, O., Lattin, J. M., & Srinivasan, V. (2008). A hidden Markov model of customer relationship dynamics. *Marketing Science*, 27(2), 185–204.
- Nitzan, I., & Libai, B. (2011). Social effects on customer retention. *Journal of Marketing*, 75(6), 24–38.
- Packard, G., & Berger, J. (2017). How language shapes word of mouth's impact. *Journal of Marketing Research*, 54(4), 572–588.
- Packard, G., & Berger, J. (2020). Thinking of you: How second-person pronouns shape cultural success. *Psychological Science*, 31(4), 397–407.
- Packard, G., Gershoff, A. D., & Wooten, D. B. (2016). When boastful word of mouth helps versus hurts social perceptions and persuasion. *Journal of Consumer Research*, 43(1), 26–43.
- Perarnau, G., Van De Weijer, J., Raducanu, B., & Álvarez, J. M. (2016). Invertible conditional GANs for image editing. arXiv preprint arXiv:1611.06355.
- Pieters, R., & Wedel, M. (2004). Attention capture and transfer in advertising: Brand, pictorial, and text-size effects. *Journal of Marketing*, 68(2), 36–50.
- Pieters, R., Wedel, M., & Batra, R. (2010). The stopping power of advertising: Measures and effects of visual complexity. *Journal of Marketing*, 74(5), 48–60.
- Platzer, M., & Reutterer, T. (2016). Ticking away the moments: Timing regularity helps to better predict customer activity. *Marketing Science*, 35(5), 779–799.
- Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. *OpenAI*, 1(8), 9.
- Rajaram, P., & Manchanda, P. (2020). Video influencers: Unboxing the mystique. arXiv preprint arXiv:2012.12311.
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., . . . Sutskever, I. (2021). Zero-shot text-to-image generation. arXiv preprint arXiv:2102.12092.
- Reisenbichler, M., Reutterer, T., Schweidel, D. A., & Daniel, D. (2022). Frontiers: Supporting content marketing with natural language generation. *Marketing Science*. doi:10.1287/mksc.2022.1354
- Rocklage, M. D., Rucker, D. D., & Nordgren, L. F. (2021). Mass-scale emotionality reveals human behaviour and marketplace success. *Nature Human Behaviour*, 1–7.
- Rossi, P. E., McCulloch, R. E., & Allenby, G. M. (1996). The value of purchase history data in target marketing. *Marketing Science*, 15(4), 321–340.
- Rust, R. T., Lemon, K. N., & Zeithaml, V. A. (2004). Return on marketing: Using customer equity to focus marketing strategy. *Journal of Marketing*, 68(1), 109–127.
- Rutz, O. J., Bucklin, R. E., & Sonnier, G. P. (2012). A latent instrumental variables approach to modeling keyword conversion in paid search advertising. *Journal of Marketing Research*, 49(3), 306–319.
- Rutz, O. J., Sonnier, G. P., & Trusov, M. (2017). A new method to aid copy testing of paid search text advertisements. *Journal of Marketing Research*, 54(6), 885–900.
- Sarkar, M., & De Bruyn, A. (2021). LSTM response models for direct marketing analytics: Replacing feature engineering with deep learning. *Journal of Interactive Marketing*, 53, 80–95.
- Schmittlein, D. C., Morrison, D. G., & Colombo, R. (1987). Counting your customers: Who-are they and what will they do next? *Management Science*, 33(1), 1–24.
- Schwartz, E. M., Bradlow, E. T., & Fader, P. S. (2014). Model selection using database characteristics: Developing a classification tree for longitudinal incidence data. *Marketing Science*, 33(2), 188–205.
- Schwartz, E. M., Bradlow, E. T., & Fader, P. S. (2017). Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science*, 36(4), 500–522.

- Schweidel, D. A., Bradlow, E. T., & Fader, P. S. (2011). Portfolio dynamics for customers of a multiservice provider. *Management Science*, 57(3), 471–486.
- Schweidel, D. A., & Fader, P. S. (2009). Dynamic changepoints revisited: An evolving process model of new product sales. *International Journal of Research in Marketing*, 26(2), 119–124.
- Schweidel, D. A., & Knox, G. (2013). Incorporating direct marketing activity into latent attrition models. *Marketing Science*, 32(3), 471–487.
- Schweidel, D. A., Yakov Bart, J., Inman, J., Stephen, A. T., Libai, B., Andrews, M., . . . Thomaz, F. (2022, forthcoming). How consumer digital signals are reshaping the customer journey. *Journal of the Academy of Marketing Science*.
- Tellis, G. J., MacInnis, D. J., Tirunillai, S., & Zhang, Y. (2019). What drives virality (sharing) of online digital content? The critical role of information, emotion, and brand prominence. *Journal of Marketing*, 83(4), 1–20.
- Timoshenko, A., & Hauser, J. R. (2019). Identifying customer needs from user-generated content. *Marketing Science*, 38(1), 1–20.
- Urban, G. L., Liberali, G., MacDonald, E., Bordley, R., & Hauser, J. R. (2014). Morphing banner advertising. *Marketing Science*, 33(1), 27–46.
- Valendin, J., Reutterer, T., Platzer, M., & Kalcher, K. (2022). Customer base analysis with recurrent neural networks. *International Journal of Research in Marketing*. forthcoming.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., . . . Polosukhin, I. (2017). Attention is all you need. In *31st conference on neural information processing systems*, 2017 (pp. 1–15). NIPS.
- Wang, Z., & Kim, H. G. (2017). Can social media marketing improve customer relationship capabilities and firm performance? Dynamic capability perspective. *Journal of Interactive Marketing*, 39, 15–26.
- Wu, G., Masia, B., Jarabo, A., Zhang, Y., Wang, L., Dai, Q., . . . Liu, Y. (2017). Light field image processing: An overview. *IEEE Journal of Selected Topics in Signal Processing*, 11(7), 926–954.
- Yang, Y., Zhang, K., & Kannan, P. K. (2021). Identifying market structure: A deep network representation learning of social engagement. *Journal of Marketing*. doi:10.1177/00222429211033585
- Yao, S., & Mela, C. F. (2011). A dynamic model of sponsored search advertising. *Marketing Science*, 30(3), 447–468.
- Zhang, Y., Bradlow, E. T., & Small, D. S. (2015). Predicting customer value using clumpiness: From RFM to RFMC. *Marketing Science*, 34(2), 195–208.
- Zhang, X., Kumar, V., & Cosguner, K. (2017). Dynamically managing a profitable email marketing program. *Journal of Marketing Research*, 54(6), 851–866.
- Zhang, S., Lee, D., Singh, P. V., & Srinivasan, K. (2021). What makes a good image? Airbnb demand analytics leveraging interpretable image features. *Management Science*.
- Zhang, Y., Moe, W. W., & Schweidel, D. A. (2017). Modeling the role of message content and influencers in social media rebroadcasting. *International Journal of Research in Marketing*, 34(1), 100–119.
- Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., & Metaxas, D. N. (2017). StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision* (pp. 5907–5915).
- Zhu, J. Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision* (pp. 2223–2232).