Schmalenbach Journal of Business Research (2025) 77:497–525
https://doi.org/10.1007/s41471-025-00219-4

ORIGINAL ARTICLE

Mining Consumer Mindset Metrics With User-Generated Content

Raoul V. Kübler · Susanne J. Adler · Lina Welke · Marko Sarstedt · Koen Pauwels

Received: 8 October 2024 / Accepted: 16 July 2025 / Published online: 4 September 2025

Abstract In the wake of digital transformation, marketers gained access to large amounts of user-generated content and data in which consumers specifically mention and discuss brands, products, and services. This data offers rich information potential and may ultimately provide marketers with the ability to use this data pool to approximate survey-based consumer mindset metrics that mirror consumer attitudes alongside the different levels of the decision-making process. We argue that leveraging this potential may ultimately help marketers overcome common limitations of survey-based metrics and enable companies to observe and track mindset metrics that have been so far inaccessible due to financial and other constraints. To this end, we propose a four-step process that first identifies the key aspects of a mindset metric based on the existing body of developed constructs, then pinpoints potential data sources, and subsequently chooses an adequate data transformation tool.

Keywords Mindset metrics · User-generated content · Construct validity · Machine learning · Customer decision journey

✉ Raoul V. Kübler · Lina Welke
ESSEC Business School, 3 Avenue Hirsch, 95800 Cergy, France
E-Mail: kubler@essec.edu

Susanne J. Adler · Marko Sarstedt
Institute for Marketing, LMU Munich School of Management, Ludwig Maximilian University of Munich, Ludwigstr. 28RG, 80539 Munich, Germany

Marko Sarstedt
Faculty of Economics and Business Administration, Babeș-Bolyai-University, Strada Teodor Mihali 58–60, 400591 Cluj-Napoca, Romania

Koen Pauwels
D'Amore-McKim School of Business, Northeastern University Boston, 450 Dodge Hall, 360 Huntington Avenue, Boston, Massachusetts 02115, USA

1 Introduction

The digital transformation of our world led to substantial changes in how much data is available to decision-makers within the last two decades. With millions of digital devices observing and tracking our everyday lives, and consumers actively creating data points whenever they browse, shop, or interact with each other, we have entered the data-richest period in human history. Today, humans and the devices they use willingly and unwillingly create more than 402.74 million terabytes of data each day (Duarte 2024), surpassing the cumulated amount of data humankind has created since the dawn of civilization. A large proportion of this data is actively created by consumers through user-generated content (UGC) such as (but not exclusively) social media posts, reviews, videos, messages, online search patterns, online service interactions, website visits, or podcasts. Users provide valuable and information-rich data with a specific focus on consumption patterns, product preferences, satisfaction information, or company and product evaluations. In light of the growing availability of data, new methods have emerged to extract information and insights from large datasets available to decision-makers (Kübler et al. 2017). Within the last decade, substantial developments in machine learning techniques, coupled with the availability of more powerful hardware, enabled decision-makers to accurately describe, predict, explain, and re-generate human behavior. Deep learning algorithms have cumulated in large language models, which, through their chat interfaces or simple application descriptions, now allow access to powerful algorithms for generating consumer insights (e.g., through the generation of synthetic datasets; Sarstedt et al. 2024a) without the need for sophisticated coding expertise (e.g., Hartmann et al. 2023). The combination of large amounts of data, machine learning methods, and powerful hardware has also enabled managers with the ability of accurately measuring, tracking, or predicting consumer behavior such as brand and product consideration and perception (Ringel and Skiera 2016), as well as ad clicks (Wang et al. 2018), purchase (Chevalier and Mayzlin 2006), or churn behavior (Khodabandehlou and Rahman 2017).

While machine learning applications have been shown to be able to ultimately predict final decision outcomes with the help of UGC, the prediction or approximation of latent variables such as consumer mindset metrics (CMMs) with UGC, has not yet been fully explored (Hair and Sarstedt 2021). Instead, managers commonly still rely on primary data to measure and track pre-purchase consumer attitudes such as brand awareness, brand consideration, or purchase intention, as well as post-purchase attitudes such as brand satisfaction, or recommendation behavior. This is especially surprising as primary data collected through surveys with consumer panels has been shown to face several restrictions such as high costs, reduced timeliness, and proneness to sampling, as well as response errors (Hulland et al. 2018). In addition, managers often rely on latent CMMs as key performance indicators in marketing, which are often included in corporate dashboards and used to evaluate the short- and long-term impact of marketing actions.

Similar to other tasks in consumer (Ringel and Skiera 2016) and market observations (Matthe et al. 2023), UGC may be an alternative data source to approximate, measure, and track CMMs. Doing so would offer multiple advantages: First, secondary data such as UGC is continuously available and can thus also be used on a daily base to monitor brand health or measure and track the impact of marketing activities. Second, UGC is freely available and is commonly not prone to traditional biases known from primary data research such as strategic answering behavior, survey fatigue, and non-response bias. Third, UGC is commonly tied to true customer experiences (i.e., product reviews or customer complaints on X) and can thus be understood as a form of implicit, revealed preferences, whereas CMMs are commonly understood as stated preferences. Finally, being implicit feedback, UGC can be considered as a more natural form of feedback than survey-based CMMs, which are—being an explicit form of feedback—known to be more prone to, for example, social desirability bias. Despite the potential of UGC for consumer mindset tracking, so far little insights exist into the suitability and how to process, use, and leverage UGC to track latent CMMs.

Addressing this gap in research, we set out to investigate and discuss the potential pros and cons of such an approach, as well as the necessary steps to facilitate the use of UGC in the context of CMMs. Our goal is to identify, systematize, and map the necessary tasks, develop guidelines for the usage of UGC to measure CMMs, and identify future research needs to ultimately develop a structured research agenda in this domain.

The remainder of this paper is structured as follows: First, we highlight and discuss the *current applications of CMMs* in marketing research and marketing practice to determine the strengths and weaknesses of the current approaches as well as to determine the needs and requirements that need to be met by a UGC-based measurement approach. Subsequently, we highlight *current applications of UGC* in modern marketing research to determine the potential of UGC for CMM measurement and potential boundary conditions. Linking insights from both sections we then *develop a framework* for combining both worlds in which we highlight necessary steps that need to be addressed to leverage the potential of UGC for CMM measurement. Finally, we develop a *research agenda* to stimulate further work in the field that is necessary to ultimately reach the goal of making UGC available to managers for measuring and tracking CMMs.

2 Consumer Mindset Metrics

2.1 Concepts of Customer Mindset Metrics

Both psychology and marketing literature identify consumer attitudes as constructs that indicate how consumers think about (cognition), feel about (affect), and act toward (conation) a brand (Vakratsas and Ambler 1999) and other objects of interest. Starting in the early 1960s (Colley 1961; Lavidge and Steiner 1961), marketing developed measures of consumer attitudes to evaluate the impact of marketing campaigns and to predict their sales effect. In their theory of buying behavior, Howard and Sheth (1969, p. 14) noted, *"Attitude is an input into executive decisions because many marketing decisions, including advertising, can be more adequately evaluated or measured in terms of attitude than of purchase behavior."* Corresponding attitude-related measures are also referred to as CMMs in that they seek to capture the *why*-question necessary to interpret observed consumer behavior. For instance, consumers may visit a brand's website because they are considering buying its products, are browsing for fun, have questions about products already bought, or want to rationalize their choice for a competing brand (Dotson et al. 2017). Pauwels and van Ewijk (2020) show that attitude-related metrics offer a superior prediction of brand sales over several months, as compared with aggregate online behavioral metrics such as the amount of weekly clicks and visits.

The literature has brought forward a multitude of CMMs. The most prominent CMMs are arguably brand awareness, brand consideration, purchase intention, brand satisfaction, brand recommendation, and brand equity. Brand awareness captures whether a brand is top-of-mind and can be measured as a consumer's ability to recall a brand's name (unaided) or as a consumer's ability to recognize a brand among other brands (aided; see, e.g., Mecredy et al. 2022; Pauwels and van Ewijk 2020). Brand consideration describes which brands consumers would regard as suitable options for usage or purchase while purchase intention captures which brands consumers would actually be willing to buy (e.g., Kübler et al. 2020; Mecredy et al. 2022). CMMs that can only be formed after having purchased and used a product or service include brand satisfaction and brand recommendation. Brand satisfaction captures whether a brand has met a consumer's expectation while brand recommendation measures a consumer's willingness to suggest the brand to others, respectively to tell others to avoid a brand (e.g., Anselmsson and Bondesson 2015; Kübler et al. 2020). Lastly, brand equity captures a consumer's preference for one brand over other brands (Washburn and Plank 2002).

2.2 The Psychometrics of Customer Mindset Metrics

As the concepts underlying CMMs are inherently unobservable, their measurement typically relies on consumers' answers to sets of survey items. These items are supposed to capture the respondents' assessment of specific traits that underlie the concept under consideration (e.g., Anselmsson and Bondesson 2015; Pauwels and van Ewijk 2020). For example, to measure customer satisfaction, researchers routinely rely on three survey items that gauge the respondents' overall satisfaction, expectancy confirmation, and perceived performance versus the customer's ideal product or service in the category (Fornell et al. 1996). From a measurement-theoretic perspective, these items are assumed to be reflections or consequences of customer satisfaction (hence, reflective measurement), thereby acting as empirical realizations of the unobserved concept; their correlation is assumed to be "caused" by the underlying concept (Sarstedt et al. 2016a). A formative measurement of customer satisfaction, on the other hand, would capture contributing traits such as satisfaction with the service, price, or product; assuming that a composite of such individual traits captures the target construct.

The logic underlying formative measurement acknowledges the approximative nature of measurement whereas reflective measures implicitly assume that any measurement can, in principle, capture a concept in full (Rigdon and Sarstedt 2022). The manifold uncertainties that come with any measurement, for example, with regard to the conceptualization, item wordings, measurement scales, data collection, and validation, however, call the assumption of perfect measurement into question (Rigdon et al. 2019, 2020; Sarstedt et al. 2024b).

Considering the practical challenges that come with collecting data on multi-item measures (e.g., with regard to careless response behavior or non-response), researchers and particularly practitioners frequently rely on single-item measures of a concept. These single items can be seen as a global assessment of the concept, such as when asking a respondent: "Overall, how do you rate your satisfaction with the company's service?" The practical advantages, however, come at the expense of reduced reliability and lower predictive validity, which may compromise their usefulness in research settings, potentially triggering flawed managerial recommendations (e.g., Diamantopoulos et al. 2012; Sarstedt et al. 2016b, c).

In their studies, managers typically customize the exact measurements to their brand or make changes in response to observed relationships, for example, pruning metrics that were too highly correlated amongst each other or did not predict behavior over time (e.g., Pauwels and Joshi 2016). This measurement flexibility also applies to measures for CMMs that are not fully standardized. Research has witnessed a proliferation of all sorts of metrics that claim to measure essentially the same concept, although often with little chance to convert one instrument's measures into any other instrument's measures (Salzberger et al. 2016). For example, research and practice have proposed a multitude of measurement instruments for corporate reputation, which rest on the same definition of the concept but differ fundamentally in terms of their underlying conceptualizations and measurement items (Sarstedt et al. 2013). Similarly, Bergkvist and Langner (2017, 2019) find considerable heterogeneity in the operationalizations of common advertising constructs, such as attitude toward the ad, attitude toward the brand, ad credibility, ad irritation, and brand purchase intention. In addition, construct conceptualizations and operationalizations change over time (Bergkvist and Eisend 2021), while the theoretical entity of interest (i.e., the conceptual variable) generally remains the same. These findings suggest there is no set way to perfectly measure a concept, even though some standardized measurement approaches can be observed. For example, brand awareness is measured as the unaided recall and aided recognition of a brand (e.g., Mecredy et al. 2022; Pauwels and van Ewijk 2020). Other CMM measures exhibit a lower consensus. To measure brand consideration, purchase intention, brand satisfaction, or recommendation intentions, consumers may select relevant brands from a list of multiple brands (e.g., Colicev and de Bruyn 2023; Kübler et al. 2020; Mecredy et al. 2022) or indicate their assessment on a Likert-scale (e.g., Anselmsson and Bondesson 2015; Cain 2022; Pauwels and van Ewijk 2020). Researchers have also suggested the use of a Net Promoter Score for which respondents indicate their willingness to recommend a brand on a scale from 0 to 10 and are sorted into endorsement categories because of their rating (Baehre et al. 2022). Measures of satisfaction range from asking participants to select brands they are satisfied, respectively not satisfied, with (Kübler et al. 2020) to using multi-item scales (Anselmsson and Bondesson 2015; Petersen et al. 2018). Moreover, the relationship between survey-measured CMMs and consumer behavior is not strong pointing toward potentially limited predictive validity (Pauwels and van Ewijk 2020).

Table 1 Popular UGC Channel Types and Examples

Channel Type	Examples	Exemplary studies
Social media platforms	Facebook, Instagram, X, TikTok, LinkedIn, Sina Weibo, Snapchat	Liu et al. (2016): Twitter (now X) Tweets to predict TV show demand
Online review platforms	Yelp, TripAdvisor, Trustpilot	Rocklage et al. (2021): Emotionality of Yelp reviews to predict restaurant table reservations
E-commerce platforms	Amazon, eBay, Etsy	Chevalier and Mayzlin (2006) use reviews to predict sales
Search engine logs	Google, Bing, Yahoo	Ringel and Skiera (2016) use Google search term volume to map consumer considerations
Blogs and microblogs	Medium, Substack, Blogger, Tumblr	Onishi and Manchanda (2012) examine impact of blog posts on brand awareness and sales
Q&A forums	Quora, Stack Overflow, Reddit	Burtch et al. (2022) examine the impact of badges on UGC creativity
Streaming platforms	Twitch, YouTube Live, Douyin, Huya Live, RedNote	Lin et al. (2021) examine impact of emotions in live streams on tips and user activity
Podcasts	Spotify, Apple Podcasts, Sound-Cloud	Kozinets et al. (2010) investigate how podcast-based seeding influences spread of UGC and eWoM

3 User-Generated Content

UGC has been in the spotlight of marketing research now for over two decades. Initially, often generally referred to as (online or electronic) word of mouth (Godes and Mayzlin 2004; Hennig-Thurau et al. 2004), the understanding and definition of UGC have become more granular over time. While the traditional digital or electronic word of mouth literature limits itself to UGC that specifically focuses on consumption-related content (Babić Rosario et al. 2020), the classic UGC literature goes beyond and includes any sort of digital information created by a private (in contrast to a paid) user's communication and online behavior and that is shared with the general public via any sort of digital (web) channel (Daugherty et al. 2008). UGC *data sources* investigated in the literature include a broad range of different channels. Table 1 provides an overview of popular channel types, along with examples. With the constant emergence of new channels, the volume and variety of digital traces left by consumers continue to expand.

A key characteristic of UGC is that it rarely occurs in a structured (i.e., numeric) *data format* (de Haan et al. 2024). If so, often simple count measures are used as approximations, such as popularity (e.g., the number of followers), brand liking (e.g., through the count of daily likes of brand posts), and brand-consumer relations (e.g., through the count of daily brand post shares). These simple approaches, however, neglect the richness of information within UGC that arises from the fact that users do not only create information by engaging with brand content, but also actively create posts, comments, videos, chats, etc. More importantly, as evidenced in the list of data sources above, the vast majority of UGC comes in unstructured (i.e., non-numeric) data such as images, text, and videos. Several studies have made use of such data to generate consumer insights. For example, the majority of research relying on UGC to approximate common brand equity metrics such as brand knowledge or brand strength have commonly drawn on textual data obtained from various online and social media sources (see, e.g., Colicev et al. 2018). Fewer studies, however, have gone one step further and used non-numeric UGC, such as images (Dzyabura et al. 2023; Hartmann et al. 2021), audio (Wang et al. 2021), or video (Zhou et al. 2021) to capture consumer or brand insights.

Unstructured data brings the advantage of richer information and thus bares more potential for generating unique insights (de Haan et al. 2024). At the same time, however, UGC data needs to be transformed into structured data in order to be further processed and analyzed. Thus, the need for *data transformation* must be considered as another key characteristic of UGC. De Haan et al. (2024) provide a detailed and structured framework for how to turn unstructured data into a suitable structured format, taking various application and company-specific contingencies into consideration.

The most commonly used tools for data transformation—at least in the context of brand measurement (Berger et al. 2020)—are sentiment measures. For textual data, sentiment measurement tools can be split into top-down and bottom-up approaches (Humphreys and Wang 2018). Top-down approaches such as the often-used LIWC or VADER commonly rely on word lists curated by linguistic research to capture sentiment or valence information. By contrast, bottom-up approaches use machine learning tools and pre-coded training data to predict the sentiment of a document by looking at the embedded words or word combinations within the document (Kübler et al. 2020). In the case of audio, either transcriptions of the spoken content are used to rely on standard textual sentiment classifiers (Zhou et al. 2021), or tools to detect loudness or tone within voice to capture sentiment (Wang et al. 2021). For images, research so far most often relies on facial expressions to capture sentiment and valence (Toisoul et al. 2021), which, however, requires images to display faces. Similarly, in the case of videos, the video material is cut into individual frames, which are then subsequently treated like images (e.g., Li et al. 2019). Besides the binary measurement of sentiment (positive and negative), an increasing number of studies now attempt to capture not only the valence, but subdimensions of emotions such as anger, fear, anxiety, sadness, joy, surprise, and anticipation (e.g., Holiday et al. 2023; Schwenzow et al. 2021; Zhou et al. 2021).

4 Connecting Conceptual Facets With UGC Insights

CMMs are typically measured with sets of items that cover multiple facets and aspects of the concept under consideration. The latent concept of satisfaction may, for example, be measured formatively by asking consumers to rate various aspects of a service (e.g., cleanliness, proximity, calmness, taste), which jointly form satisfaction. In case of a reflective approach, the same concept would be measured by the attempt to capture the outcome of the construct such as, "How satisfied are you with the brand?" and "How well does the brand match your expectations?" (Fornell et al. 1996). Traditional CMMs measurement therefore captures stated preferences, collected through explicit feedback. In contrast, UGC inherently reflects aspects of revealed preferences embedded in implicit feedback. This content often contains

Table 2 Advantages and Disadvantages of Traditional Data versus UGC for CMM Measurements

	Advantages	Disadvantages
Traditional data	*Sampling:* Potential to ensure and check for the representativeness of a sample for a specific population Targeted sampling of, for example, specific groups	*Sampling:* Potentially high data collection costs for recruiting and incentivizing study participants Reduced timeliness Prone to sampling errors including non-response biases Potential careless response behavior
	Scale use: Direct measurements, for example, concerning attitudes Validated, standardized measures available	*Scale use:* Scale-specific disadvantages (e.g., reduced reliability and lower predictive validity in single-item scales) Many implicit assumptions in scale design increase uncertainty (e.g., scale types, wording of items and answer categories) Use of unstandardized, ad-hoc measures Scale modifications may change the nature of measurement; necessary scope often unclear Emergence of response biases (e.g., social desirability bias)
UGC	Captures real-world consumer behavior Available from multiple sources, continuously available, and often freely available Rich information from combining multiple data sources available Less prone to traditional biases known from primary data research (e.g., social desirability, survey fatigue, non-response bias)	Uncertain data composition in terms of target groups who provide UGC Self-selection bias (e.g., dominance of negative responses and data from very active users) Risk of incomplete construct assessment Measures are mostly indirect and approximative Object of measurement may be ambiguous (e.g., product-related vs. brand-related assessments) Combination of data from various sources can be prohibitively complex and expensive Data is mostly unstructured, and transformation is often necessary

information depicted by the items that operationalize CMM constructs. Consumers may, for example, compare products in social media videos (e.g., product haul videos) or other types of social media posts by revealing which brands they consider and which brand they prefer most. Similarly, consumers may speak in reviews about aspects they like or dislike with a product or how satisfied they are with a product or service to ultimately even give a recommendation to other users. However, by simply counting occurrences or only assessing the sentiment of a post or review, researchers risk missing relevant construct facets, which may bias the validity and reliability of the obtained measure. While we may grasp the tone of a post, we fail to capture its content or the ultimate meaning of a UGC element, both of which are essential for adequately capturing the construct.

Overall, UGC as well as scale measures for CMM come with distinct advantages and disadvantages (Table 2). To enable UGC to capture the aspects of each CMM in a valid (and reliable) way, we suggest combining the information provided by the extensive body of psychometric research with the information potential of UGC together with the abilities of today's rich machine learning capabilities to capture the necessary nuances unique to each construct. In the following, we develop a conceptual guide that derives the necessary steps to subsequently discuss each step in more detail.

5 Marrying Psychometrics With UGC: A Four-Step Process

To embed UGC in a psychometric framework, we suggest the four-step process shown in Fig. 1. The process starts with the examination of the properties, characteristics, and aspects of a CMM based on the items used by the available, theoretically developed, and empirically tested scales (Step 1). Researchers should distinguish between brand and product level specific metrics; for example, metrics such as awareness, consideration or satisfaction may either be brand (macro) or product (micro)-related. Depending on whether a micro- or macro-perspective is taken, dif-

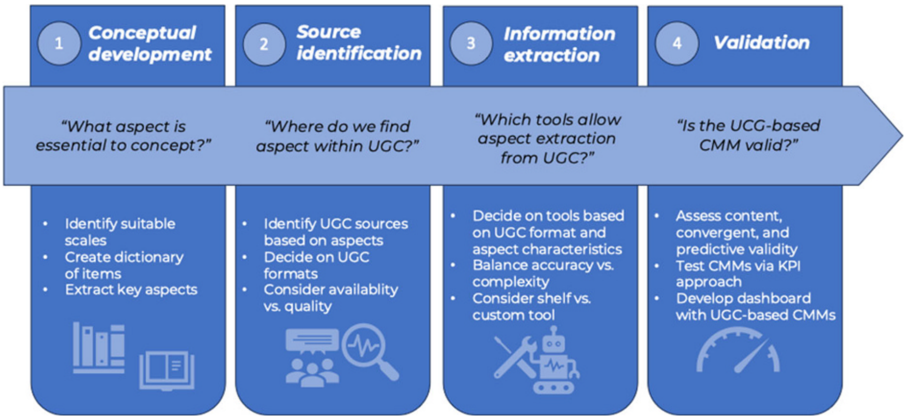

Fig. 1 Four Step UGC Mindset Metric Measurement Process

ferent data sources and aspects are used. For example, Ringel and Skiera (2016) rely on product-related search terms to identify alternative products within the consideration set and Matthé et al. (2023) rely on investor relation reports to map brand evolution over time. Similarly, measuring brand level effects may require researchers to ensure that their UGC identification process captures all products and services of a company, while on product level one limited set of product specific anchors may be sufficient. By examining the items used to measure the CMM, researchers can identify the key aspects tied to the construct, which will subsequently guide the UGC identification process.

In Step 2, the type of CMM and its corresponding aspects will determine the selection of UGC sources relied upon in the data collection process. Source selection will, on the one hand, depend on the information consumers share in specific channels and how well that information can be used to measure the aspects of interest. Secondary factors such as accessibility, available data volume, costs, and legal aspects may also guide the source selection process. We suggest that these secondary aspects become more important in case researchers attempt to continuously collect data instead of using data sources only once or occasionally. The choice of a data source is also going to influence the format of the data received (i.e., structured vs. unstructured) and will thus also affect the decisions in Step 3 in which the researchers identify a suitable aspect information extraction tool.

In this third step, the focus is on determining how to extract the identified latent aspects identified in Step 1 from the data obtained in Step 2. The method selection will largely be guided by the previously identified metric aspects and the data characteristics, but also by secondary factors such as time, costs, and the question if insights are needed on a continuous basis (e.g., in a dashboard) or at a few selected time points (e.g., before and after a campaign).

Finally, given the importance of CMMs for tracking consumer attitudes along the decision-making funnel one needs to control if the obtained measures meet quality standards. We therefore suggest that in Step 4, researchers assess the measure's validity, focusing particularly on the predictive power of the obtained UGC metrics.

In the following, we discuss potential approaches for each step with the aim to encourage future research in the domain to rely on the suggested sources and methods and to add empirical evidence to our suggested four step model.

5.1 Conceptual Development

Given the extensive stream of scale-related research, we do not intend to repeat a description of scale-related aspects (Step 1), but refer to the relevant literature (e.g., Bruner 2023; DeVellis and Thorpe 2021; Jebb et al. 2021). Instead, our elaborations will focus on the questions of where to find data (Step 2), how to extract information (Step 3), and how to ensure that the resulting measures are reliable and valid (Step 4).

5.2 Source Identification

Previous research has relied on various UGC sources to explore the relationship between UGC and various CMMs along the decision-making funnel, as well as related outcome measures. Commonly studied sources include social media platforms (e.g., Colicev et al. 2018; Kübler et al. 2020; Liu et al. 2016), online review and e-commerce platforms (e.g., Chakravarty et al. 2010; Chevalier and Mayzlin 2006; Rocklage et al. 2021), and blogs (e.g., Gopinath et al. 2013; Onishi and Manchanda 2012). For example, Kübler et al. (2020) use Facebook data to evaluate the effectiveness of sentiment measurement approaches in capturing daily CMMs. Similarly, Colicev et al. (2018) link owned and earned social media measures to brand awareness, purchase intent, and customer satisfaction. In contrast, other studies focus on the impact of UGC on key outcome measures such as sales or revenue, thereby indirectly underscoring its potential to approximate CMMs. For instance, Chevalier and Mayzlin (2006) link the volume and valence of online ratings to the sales rank of books on Amazon, while Rocklage et al. (2021) use consumer reviews to predict movie box office revenues, book sales, new brand followers, and restaurant reservations.

While these findings highlight the significant influence of UGC on consumer behavior across various contexts, research exploring links between CMMs and other UGC-rich platforms, such as streaming or podcast services, remains limited. These platforms, however, offer unique UGC types, such as user interactions, content

More structured data formats ↔ ***More unstructured data formats***

	Brand awareness	Brand consideration	Purchase intention	Brand satisfaction	Brand recommendation	Brand equity
Social media platforms	✓✓✓	✓✓	✓	✓✓	✓✓	✓✓
Online review platforms	✓✓	✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓
E-commerce platforms	✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓
Search engine logs	✓✓	✓✓✓	✓✓✓	✓	✓	✓
Blogs and microblogs	✓	✓✓	✓✓	✓✓	✓✓✓	✓✓
Q&A forums	✓	✓✓	✓✓✓	✓✓	✓✓	✓
Streaming platforms	✓✓	✓	✓	✓✓	✓✓✓	✓
Podcast platforms	✓	✓	✓	✓✓	✓✓✓	✓

✓✓✓: Strong suitability; ✓✓: Moderate suitability; ✓: Limited suitability

Fig. 2 UGC Data Sources

consumption patterns, and engagement metrics, which could yield valuable insights into consumer behavior and decision-making processes. As such, they also represent promising data sources for approximating CMMs through UGC. Fig. 2 provides a summary of potential UGC data sources and indicates the suitability of each source for obtaining information relevant to a CMM[1].

The suitability of a data source for measuring a specific CMM largely depends on whether it includes structured and/or unstructured data formats. For instance, structured metrics such as the number of views and shares or the frequency of brand mentions and hashtags may be useful for gauging brand awareness as they reflect how frequently a brand is seen or discussed, serving as a proxy for its visibility and reach. In contrast, these metrics may not effectively capture more nuanced consumer attitudes such as brand recommendation which often require a more in-depth analysis of unstructured data. For example, brand recommendation is indicated by data that reveals advocacy or a willingness to promote the brand to others such as comments or reviews explicitly recommending the brand. In support of this distinction, Liu et al. (2016) find that the pure volume of tweets or Google searches 24 h before a TV show aired are poor predictors of subsequent ratings (i.e., recommendation). Conversely, tweet content proved to be a much stronger predictor of TV ratings.

Given the varying importance of structured and unstructured data in assessing different CMMs, it is essential to consider which data sources contain which data format(s) to determine their suitability. Social media platforms, for instance, offer structured metrics such as views, likes, shares, and comments, alongside unstructured data like the content of posts and written comments. Similarly, online review platforms feature structured measures, including product star ratings and review helpfulness votes, as well as unstructured data such as written reviews and user demographics. Consequently, both data sources are well-suited for measuring CMMs at different stages of the marketing funnel, from awareness to post-purchase evaluation. However, as social media platforms typically contain a greater variety of structured data (i.e., number of views, likes, shares, comments) than online review platforms, we consider them as even better suited to measure brand awareness. Conversely, as online review platforms offer more detailed written information, they are generally better equipped to measure brand recommendation. Evidence for this differentiation comes from Kübler et al. (2020) who demonstrate that volume metrics from Facebook are more effective in explaining brand awareness than purchase intent, whereas a sentiment classifier trained on Amazon reviews better predicts brand satisfaction and recommendation. Similarly, Rocklage et al. (2021) show that whereas star ratings on online review platforms are poor predictors of product success, the emotionality expressed in reviews is a much stronger predictor.

In contrast, search engine logs primarily feature structured data in the form of search volumes and trends, such as those available through platforms like Google Trends. This data reflects how often a brand is being sought after, serving as a reliable indicator of its visibility and popularity in the market. However, Google Trends lacks

[1] Guyt et al. (2024) and Yildirim and Kübler (2023) provide detailed guidance on extracting UGC from web data sources using web crawlers or APIs. We thus refrain from technical explanations and refer to these sources.

Table 3 UGC Information Extraction Approaches

Metric	Characteristics	Sample Items	Text-based sources	Image-based sources	Video-based sources
Brand awareness	The ability of a brand to shine out of clutter, being separable, to distinguish itself from the competition	I can recognize X among other competing brands. I am aware of X. Some characteristics of X come to my mind quickly. I can quickly recall the symbol or logo of X. I have difficulty in imagining X in my mind.	Use simple *keyword matching* for direct mentions and hashtags associated with the brand to track the frequency of brand mentions, capturing general popularity. Rely on the *number of followers*. Identify the number of mentions with other brands by, e.g., using a *topic model* to track competing brands being mentioned alongside your brand to build a share of voice score. Rely on NLP tools such as *rapid automatic keyword extraction, Word2Vec, or dependency parsing* to identify common terms associated with the brand, and understand the dominant brand characteristics.	Use *image/logo recognition* to detect the focal brand's product, logo, or packaging in images to capture general popularity. Use *image/logo recognition* to detect focal brand's and category competitors' products, logos, or packaging in images to capture share of voice. Use *object classifiers* to identify common objects and context factors often appearing with focal brands to understand use context.	Use text tools to analyze *captions, voiceovers, or comments* to derive an understanding of brand and competitor mentions. Use *logo and symbol detection models* to identify *how frequently and how long* the focal brand's logo or symbols appear in videos. Rely on *product haul* videos and object identifiers to understand brand popularity in comparison to other brands and to capture the share of voice.
Brand consideration	The likelihood that consumers regard a brand as a relevant option for a purchase	Which of these brands would you consider buying?	Use simple *keyword matching* for elements like wish lists, active comparisons, search requests, etc. Apply *embedded vector space models* to understand the co-occurrence of brand mentions in the consumption context. Use *topic models* to identify comparative language cues.	Identify *hashtags* within captions to identify posts that compare brands. Use o*bjects and logos* to create a *vector space model* based on co-occurrences.	Use *hashtags* within captions, perform *NLP text analysis* with transcripts, or analyze *comments* under video content. Identif*y objects* within video frames. Identify which *other brands* occur and how often they appear alongside the focal brand. Determine how well brands are covered in *subtitles or audio transcripts*.

Table 3 (Continued)

Metric	Characteristics	Sample Items	Text-based sources	Image-based sources	Video-based sources
Purchase intention	Consumers' willingness to buy a brand	From which of these would you be most likely to purchase? Below are some brands of product category Y. How likely would you buy Brand X in the future?	Search for *future-orientated language* and *key phrases* indicating future purchase intent, such as: – "I can't wait to buy X," – "I'm definitely getting X," – "Next time, I'm buying X instead of Y." Use *NLP tools* to detect comparative language like "better," "prefer," or "more likely to buy," to gauge which brands are more likely to be purchased. Count *engagement metrics* on purchase-related social media posts.	Apply *NLP tools* to captions and comments. Count *shopping tags and referral links*. Identify *PoS content and situations* within images.	Apply *NLP tools* to captions and audio transcripts. Identify *PoS content and situations* within video frames.
Brand satisfaction	The ability of a brand to fulfill a customer's needs and expectations	How satisfied are you with brand X? How well does brand X match your expectations? Imagine a perfect brand in this category. How close to this ideal is brand X? Of which of the following brands would you say that you are a "satisfied customer"? Of which of the following brands would you say that you are a "dissatisfied customer"?	Use *satisfaction-related dictionaries* to identify satisfaction context (e.g., expected, wanted, etc.) together with sentiment score from, e.g., VADER. Use *emotion analysis* (EmoLex, SieBERT). Apply *topic models/ABSA* which focus on the identification of product/service aspects together with sentiment scores. Extract *numeric ratings* from reviews.	Apply *NLP tools* to captions, tags, or comments. Use *computer vision tools* to identify aesthetically pleasing or celebratory contexts that suggest satisfaction, especially in cases of *unboxing or product review images*.	Apply *NLP tools* to audio transcripts, captions, tags, or comments. Rely on *facial expressions, gestures, and tone* to understand if people support the product/brand.

Table 3 (Continued)

Metric	Characteristics	Sample Items	Text-based sources	Image-based sources	Video-based sources
Brand recommendation	Consumers' willingness to encourage others to buy or use a brand	Which of the following brands would you recommend to a friend or colleague? Which of the following brands would you tell a friend or colleague to avoid? How likely is it that you would recommend brand X to a friend or colleague?	Use *word lists/dictionaries* to track recommendation-specific posts. Extract *Net Promotor Score*-like elements (e.g., 10 out of 10). Use *topic models* to identify and track testimonial statements. Extract and collect *recommendation scores* from, e.g., reviews. Identify UGC where users *tag others* recommending the brand.	Apply *NLP tools* to captions, tags, or comments. Use *computer vision tools* to identify *interactive elements* on visual social media such as interactive stories, polls, or recommender links. Use *computer vision tools* to identify the context in which *testimonials* present the product.	Apply *NLP tools* to audio transcripts, captions, tags, or comments. Apply *computer vision tools* to video frames.
Brand equity	The ability of a brand to affect buying decisions based on objective or subjective characteristics consumers value more for the focal brand than for its competitors	It makes sense to buy X instead of any other brand, even if they are the same. Even if another brand has the same features as X, I would prefer to buy X. If there is another brand as good as X, I prefer to buy X. If another brand is not different from X in any way, it seems smarter to purchase X.	Count the number of statements where users *express a rationale for choosing a focal brand over others.* Rely on *word embedding models* to identify the number of instances in which UGC acknowledges that competitors might be "as good" but still express preference for the focal brand. Rely on *comparison mining methods* (use dependency parsing or other syntactic tools to extract comparative statements involving focal brands and competitors). Use *ABSA* to conduct reasoning extraction, by employing a logic classifier that identifies phrases where users rationalize or justify their choices.	Apply *NLP tools to captions*, search for hashtags, and other identifiers. Rely on *object classifiers* to extract logos and understand which brands are more often or more prominently featured. Rely on *object classifiers* to grasp setting and apply *word vector models* to identify objects to understand the fit between content and brand. Count the number of *repeated brand features/occurrences* in a specific setting or for specific keywords or search terms.	Apply *NLP tools to captions*, search for hashtags, and other identifiers. Create *transcripts from voiceover* and apply NLP tools to obtained text. Rely on *vector-space models* of brand mentions to identify comparison across brands, rely on distance from model to understand strength.

unstructured data, which limits its effectiveness in measuring consumer attitudes like brand satisfaction or recommendation.

Podcast platforms primarily provide largely unstructured data through transcripts rather than numeric engagement metrics such as likes or stream counts. While these transcripts can offer rich insights into consumer evaluations and discussions about a brand, they are less useful for gauging brand awareness at the beginning of the marketing funnel.

5.3 Information Extraction

In the third step, researchers need to extract the aspect information from the UGC with the help of adequate tools. As shown in Table 3, various approaches and methods are suitable for this task. The choice of a method depends on the data format (structured vs. unstructured), the complexity of the data, the type of data (text vs. images vs. videos), as well as budget and time constraints.

In the case of top and bottom of the funnel metrics, such as brand awareness and satisfaction, where numeric information is available and easily accessible, simpler methods that rely on count data may be sufficient to provide initial insights (see, e.g., Kübler and Seggie 2024). Tracking the number of followers of social media platforms or counting the number of review ratings can validly depict brand awareness or satisfaction, as demonstrated by numerous studies in the field using such data (see, e.g., Hewett et al. 2016; Colicev et al. 2018; or Kübler et al. 2018) to approximate or explain brand awareness or satisfaction. However, understanding the various facets of these constructs requires a more nuanced approach that draws on unstructured data and applies more complex tools.

For CMMs, where the identification of brand mentions to capture brand awareness or brand consideration, is essential, researchers may rely on tools developed for object identification in unstructured data. This may involve brand name lists (or dictionaries) to tag social media posts or calculate share of voice scores by comparing how many times the own brand is mentioned compared to others (e.g., Ringel and Skiera 2016). For non-textual data, object classifiers or object detector models may be used to recognize brand logos or key products to achieve a similar outcome. Besides using one of the various commercial off-the-shelf solutions built into AWS and MS Azure or open-access tools hosted on platforms such as HuggingFace (see, e.g., https://huggingface.co/spaces/nathanjc/Logo_detection_YoloV7), researchers may build their own detector models. Training and fine-tuning such classifiers have become easier with various packages in R or Python, offering the application of state-of-the-art detection models—see, for example, Yildirim and Kübler (2023) for a logo identifier markdown approach. The availability of pre-trained models on platforms such as HuggingFace that only require a "few shots" of training data (commonly less than 100 images per object) to deliver reliable results, has also significantly contributed to the dissemination of classifiers—see, for example, Carion et al. (2020) for a transformer-based object detection model.

In situations where not only brand awareness needs to be tracked, but also brand-related elements such as dimensions of brand equity (e.g., Yoo and Donthu 2001), other machine learning-based techniques may be helpful. For textual data, topic models, which allow the identification of latent aspects within text data (Büschken and Allenby 2016), may help extract such insights. The key idea of these models is that specific words are tied to topics within documents and that documents are composed of different topics (Blei et al. 2003). Different forms of topic models are available with Latent Dirichlet Allocation (LDA) models being the most widely applied ones (e.g., Adler and Sarstedt 2021; Adler et al. 2024). However, researchers increasingly apply structural topic models (STM), which can account for evolving trends over time (Roberts et al. 2019), or models that account for word embeddings such as BERTopic (Grootendorst 2022). For more details and a step-by-step guide on how to estimate topic models, see Yildirim and Kübler (2023).

In the case of image (and video) data, object classifiers can similarly be used to capture contextual information and derive insights about various aspects or the meaning of a brand. By counting or measuring the types of objects commonly shown alongside a brand or product, researchers may obtain better information about how consumers perceive brands. In some cases, this may require a combination of tools, where an object detector is used to identify all objects occurring with a brand, and word-vector models are subsequently used to automatically derive meaning from the obtained brand-object-embeddings—see, for example, the approach of context fit measurement that Kübler et al. (2024) present.

As shown in Table 3, the ability to capture and integrate contextual information becomes increasingly relevant for mid-funnel CMMs, such as brand consideration and purchase intention. For the former construct, when relying on textual data, pre-curated word lists, topic models or vector models may be used to identify which brands are commonly mentioned together in a consumption context. Similarly, researchers may employ these methods to determine which other brands are frequently mentioned or searched alongside their focal brand to depict the consumers' consideration set (e.g., Ringel and Skiera 2016). Subsequently, the share of voice or mentions can be used to understand the position within the consideration set as a measure for the strength of consideration. The same can be achieved with the help of object classifiers in the context of visual information, such as images or videos.

While share of voice may suffice to capture brand consideration, purchase intention may require a more detailed understanding of how strongly a UGC element suggests that the content creator prefers a mentioned brand or product over its alternatives (see, e.g., Karniouchina et al. 2022). To address this concern, researchers may rely on textual data using either word dictionaries (Pennebaker 2001) that detect future- and goal-oriented language (like LIWC) or part-of-speech taggers that find instances where users indicate clear buying intentions (e.g., through the use of words like "can't wait to get it" or "going to buy") or preferences (e.g., through words such as "better than" or "would buy X instead of Y") as, e.g., demonstrated by Sepehri et al. (2023). Similarly, sentiment scores can be applied to UGC that discusses or compares alternatives. In particular, aspect-based sentiment analysis (ABSA) can be leveraged to assign a sentiment score to each brand mentioned in a UGC piece (Do et al. 2019). ABSA can be understood as a combination of topic modeling and sentiment analysis. Pre-trained state-of-the-art models are available on all common platforms and can be fine-tuned with few shots (see, e.g., SetFitABSA on https://huggingface.co/blog/setfit-absa). Fine-tuning requires researchers to provide the model with training data that features aspects (or in our case, brands) and examples of positive and negative brand-specific comments (Liu et al. 2020). For images and videos, one may apply the same techniques to textual data coming with an image or video, such as captions, comments, subtitles, or transcripts of voiceovers (see, e.g., Ghosh and Deb 2022). Meanwhile, the identification of purchase-specific UGC may be achievable by first trying to identify images or videos that are tied to a purchase context (e.g., point of sale content, unboxing posts, or product haul videos) to subsequently identify which brands are featured more often, as demonstrated by Orti et al. (2019).

The importance of capturing both context and content within UGC to accurately understand the relevant aspects becomes even more important for the post-purchase CMMs. In the case of satisfaction, one may directly assess satisfaction by the source information (i.e., by examining reviews or review-specific content, such as review images or product review videos) and rely on the emotions expressed within this content to capture the direction and the strength of satisfaction. Emotions can be captured in different ways. For textual data, emotion dictionaries such as EmoLex (often also referred to as NRC; Mohammad and Turney 2013) allow a quick and easy measurement of various emotional dimensions such as joy, surprise, anticipation, anger, disgust, and fear. However, EmoLex suffers from the common dictionary-based shortcoming as it does not capture negations and is only available for a few major languages (Berger et al. 2020). Hartmann et al. (2023) provide a pre-trained sentiment analysis transformer-based model that captures the same emotional dimensions, offers greater accuracy, and can be fine-tuned to brand-specific contexts. The resulting information can then be used to create a satisfaction index, as demonstrated by the online video game distribution platform Steam, which also uses an emotion classifier—accessible through Steam's main API—to derive a satisfaction score from its review texts (Guzsvinecz and Szűcs 2023). Hotz-Behofsits et al. (2025) furthermore present a new tool that first "emojifies" textual information to then infer emotions from the emojified content. Similarly, computer vision as a research discipline that attempts to transform unstructured visual content into a structured format (Voulodimos et al. 2018) provides tools for the identification of facial emotions, allowing to capture an emotion score from product review video frames.

Both approaches can also be applied to measure brand recommendation, a CMM at the bottom of the funnel. Here, one may actively use word lists or topic models to identify whether a piece of UGC contains a specific recommendation, before applying the previously discussed NLP and computer vision tools to determine which brands are recommended and for what reasons. In addition, referrer and affiliate links, poll mechanisms, and scores that resemble classic ratings or net promoter score voting elements (such as "I would give 9 out of 10" or "for me a clear five-star hotel") may be used to identify recommendation-specific content. In this case, the numeric information can also be leveraged to understand how strongly a piece recommends a brand or product.

5.4 Validation

Once one has collected and transformed the data with the help of the methods described above, the question arises, how well the obtained measures depict the CMM and if the obtained measures can be used to track customer mindset along the decision-making funnel.

The validation of UGC measures may, in principle, follow standard approaches well-known from the psychometrics literature, while accounting for the specificities of the data type. This process starts with evaluating how well the UGC measure captures the concept domain identified in Step 1 (Fig. 1). Such an evaluation is subjective, but systematic in nature, typically drawing on researchers with expert knowledge in the domain. This qualitative assessment should be followed by a quantitative analysis that examines the degree to which the UCG measure correlates with an alternative measure of the same concept (convergent validity) and its ability to predict relevant outcome variables (predictive power).

The straightforward approach to assess convergent validity is to use survey-based CMMs and estimate the correlation between the survey and UGC data (e.g., Kübler et al. 2020). Such an analysis, however, requires researchers to have control data at hand, which is often not the case due to financial constraints or because such data has not been collected before and is thus not available. In some cases, researchers may draw on secondary data proxies as surrogates for primary survey data. For example, research institutions have used customer complaints as a proxy for (dis)satisfaction (Hunt 1991). Such secondary data proxies rarely cover the concept's domain in full, but may still act as reasonable standards of comparison, provided that the literature offers support for their use (Houston 2002). Without such theoretical support, examining the proxy's correlation with other UGC measures within the study can offer clues regarding discriminant validity (i.e., the degree to which measures of different concepts are sufficiently empirically distinct).

The primary concern in the validation process should be establishing the UCG measure's predictive validity to support its relevance for managerial decision-making (e.g., Sarstedt and Danks 2022). Predictive validity assessment may also draw on survey-based control data, but in the absence of such data, researchers may test the obtained metrics by relying on the general concept of the decision-making funnel. According to this concept, investments at each level of the funnel (e.g., awareness) will, over time, move down the funnel and affect lower levels of the decision-making process (e.g., Colicev et al. 2018). Knowing that each CMM level is linked to (or predicts) a lower-level CMM and, in turn, to observable outcomes, one can develop a traditional KPI-type model. Recognizing that there is a causal chain between awareness, interest, consideration, purchase intention, satisfaction, and recommendation, researchers may attempt to tie the metrics obtained from UGC to performance variables that are observable at each stage following the lead performance indicator approach that Pauwels (2014) describes. By understanding how well each UGC-based CMM measurement predicts metrics further down the funnel, alongside performance variables available within the company (such as leads, purchases, or repeat purchases), researchers may be able to confirm the validity of the obtained UGC measures. For a guide on how to conceptually build and empirically test such a model, we refer to Hanssens et al. (2014) and Yildirim and Kübler (2023).

While establishing convergent and predictive validity should be relevant for all types of UGC measures, researchers may consider further validity types, depending on the nature of the data and its source. For example, Berger et al. (2020) provide a comprehensive framework for validating measures derived from text analysis. Rust et al. (2021) propose a real-time brand reputation tracker based on social media mining using a customer equity framework. Their work highlights the value of theory-driven UGC analysis and provides robust validation techniques by linking brand drivers to financial outcomes. Finally, Houston (2002) introduced a validation guideline for measurements of marketing constructs that draw on secondary data. This procedure comprises a three-step process, which partly overlaps with our framework (e.g., theoretical specification), but considers additional elements such as one-dimensionality and nomological validity assessment.

6 Discussion

Understanding how UGC translates into consumer perceptions, attitudes, and intentions has become crucial for managerial decision-making. While marketing researchers and practitioners often rely on numeric data (e.g., like counts) to approximate consumer assessments, UGC offers much richer information that can be reaped with today's machine learning tools. Addressing the challenges that come with such an approach, we developed a four-step process that uses item contents identified in psychometric assessments of CMMs as a blueprint for source identification as well as information extraction and control. To facilitate the adoption of our framework, we classify UGC data sources according to their suitability for capturing prominent CMMs and identify tools to do so. Applying our framework, marketing researchers and practitioners can efficiently leverage UGC data to measure CMMs. This approach offers further benefits to the field: The combination of different UCG measures in terms of, for example, numeric data (e.g., likes on social media) and text data (e.g., reviews on retailer websites) can help deepen the understanding of a CMM through triangulation. Insights from UGC data can also be combined with scale measurements from customer surveys to bring the best of the old and the new worlds together. Additionally, due to its timeliness and lower level of theoretical input, UGC can also help improve CMM's survey measurements. Specifically, comparing results from UGC to survey items indicates which scale dimensions or items are represented well in UGC but also which aspects are missing. Indeed, consumers' online reviews could help identify which product or service features are likely linked to customer satisfaction. Such assessments can inform formative measurement approaches to ensure that the items capture the concept in full. Similarly, the absence of themes in UGC could help refine scales. Respondents in surveys must often answer all items to proceed in the questionnaire—irrespective of whether a measurement is actually useful (Avis et al. 2014). If a repeated assessment on the grounds of UGC does not identify a specific measurement aspect, it may not be relevant and could be considered for removal in future psychometric assessments. UGC can therefore

Table 4 Agenda of Future Research

Research Agenda	Specific Research Questions
Conceptualization	Which aspects are suitable for product-specific CMM characteristics? Which aspects are suitable for brand-specific CMM characteristics? How strongly can UGC based CMM measures be standardized vs. need of customization? Which metric/UGC data combination produce which level of uncertainty?
Source Identification	Which data source is suitable for which CMM aspect? How to identify brand- and product relevant UGC for CMM measurement? How to sample UGC data according to the target group? Which mix of UGC sources is necessary to ensure representativeness? How to avoid J-shape trap of valence? How to best avoid data dependence and ensure permanent data access? How much is UGC data prone to manipulations and how to avoid fake insights?
Information Extraction	Which content measure tool for which CMM aspect and data source? Which CMM aspects generally require which extraction tool traits? What are costs vs. benefits trade-offs of measurement approaches in terms of CMM measure improvement? What are costs and benefits of multimodal vs. singular model measurement approaches
Validation	What are suitable validation guidelines for UGC-based CMMs? Do we need different guidelines for different types of CMMs? Which reliability and validity types should be considered in the validation of UGC-based CMMs? How can fake reviews and AI-generated content be identified?

serve as an outside validity check for established scales that would be specifically useful if the construct is context-dependent or time-variant. Due to the variety of channels that produce UGC, researchers may also draw on a wide range of data from diverse populations (e.g., teens and younger adults on TikTok and business professionals on LinkedIn). With their specific socio-demographics, experiences, and attitudes, using UGC for scale refinement may aid or replace exploratory steps in scale development.

7 Future Research

While our framework sets the stage for using UGC for CMM measurement, it also offers room for follow-up research, which we summarize in Table 4.

As for the identification of key CMM construct aspects, we first focused on extending measurements of CMMs with UGC. Compared to the unstructured data that UGC provides, CMMs come in the form of standardized measurements. However, measures validated on psychometric grounds provide an even higher level of standardization and can be accompanied by norms that identify which values on a scale are high or low compared to a specific population (Rigdon and Sarstedt 2022). This level of standardization is, as of now, not achieved by measures relying on UGC. Future research should thus prioritize the delivery of an understanding whether UGC CMM measures always need to be customized or if they can also be standardized, and if so, what is required for this process.

Relatedly, in applying our framework, researchers need to explicitly acknowledge that UGC data only allows for approximating the concepts of interest. Recent research highlights the role of measurement uncertainty in psychometric assessments of concepts in this context, which induces a validity gap between the theoretical concept and the construct (e.g., Rigdon et al. 2019). Using UGC data is likely to widen this gap due to the idiosyncrasies of the data and the researcher's degrees of freedom in the processing and analysis of the data. The use of UGC in CMM measurement therefore emphasizes Rigdon et al. (2019) call to put greater effort into quantifying measurement uncertainty (see also Rigdon and Sarstedt 2022). Future studies should also develop validation guidelines for metrics derived from UGC. While content, convergent, and predictive validity, as highlighted above, are important elements, other aspects of validity are also relevant. Corresponding studies may build on the conceptual works of Berger et al. (2020) and Rust et al. (2021), or draw on extant guidelines proposed in related contexts (e.g., secondary data validation; Houston 2002). The goal should be to develop a comprehensive approach that enables researchers to assess the validity of a given metric based on UGC traits, CMM aspects, and external validation data characteristics. This approach should help determine how well a metric aligns with existing external data sources and evaluate its validity holistically. In doing so, future research should thus also target the question of which aspects drive or minimize metrological uncertainty, which has been identified as a major contributor to low replication rates (Rigdon et al. 2019).

Regarding data source management, future research should first establish a clearer understanding of which sources—and combinations of sources—are most suitable for UGC-based CMM measurement. This involves assessing the representativeness of data sources for a given target group and, on the other, developing a more robust approach to constructing a valid sample from diverse UGC sources. A key challenge lies in the inherently polarized nature of UGC, where users tend to either strongly praise or harshly criticize a brand or product. To mitigate the well-known J-shaped distribution issue observed in online review ratings (Hydock et al. 2020), researchers must explore effective strategies for capturing neutral UGC as well. Finally, it is important to recognize that UGC is susceptible to manipulation, such as fake reviews. Future research should therefore explore methods for detecting and removing invalid or manipulative content from samples.

Relatedly, the quality of UGC-based CMM measures can be compromised by the presence of non-organic content, such as fake reviews or AI-generated texts (e.g., Kovács 2024). Fake reviews, often generated for promotional or malicious purposes, can introduce significant biases in CMM measures by distorting sentiment distributions, misleading consumer perception analyses, and reducing the overall credibility of insights derived from UGC. Despite efforts by platforms to detect and filter fraudulent content, evolving techniques in automated review generation and coordinated manipulation campaigns continue to pose challenges (Wu et al. 2020). Similarly, the rise of AI-generated texts introduces another layer of complexity. With advancements in natural language generation, AI systems can produce highly convincing UGC, ranging from product reviews to social media posts. While AI-generated content can enhance engagement and streamline communication, it also risks undermining the validity of CMM measures if such content lacks genuine consumer intent or introduces artificial sentiment trends, particularly if these are being reinforced by the reuse of AI-generated data (Xing et al. 2025). To mitigate these risks, future research and practice should focus on developing robust filtering mechanisms, integrating AI-powered detection tools, and refining validity controls for UGC-based CMM measures.

While UGC is widely available online, accessibility remains a critical challenge for researchers. Many platforms increasingly restrict automated data retrieval through their terms of service, advanced technical measures, or limited API access (Boegershausen et al. 2022), thereby seeking to maintain competitive advantages. For instance, platforms like Amazon or Google employ advanced measures to prevent web scraping. Public APIs, by contrast, offer a more efficient alternative but are often expensive, rate-limited, or subject to lengthy application processes. For example, X charges $ 5000 per month for its Pro API to retrieve 1,000,000 posts (X, 2025), while TikTok's free academic API is restricted to select regions and institutions (TikTok 2024) Such restrictions may intensify in the future, posing significant barriers to researchers and disproportionately impacting smaller or mid-sized companies that lack the resources to obtain proprietary datasets or access expensive APIs. Reflecting on these dynamics, future research should address the implications of limited data accessibility and explore alternative ways to gather and utilize UGC data.

As de Haan et al. (2024) and Kübler (2023) highlight, research has provided a rich and expanding array of UGC information extraction tools. New methods and tools often lead to classic "benchmarking" studies, which compare the suitability of a new tool against existing tools. This chain of research becomes increasingly difficult to follow due to the rapid development of methods. We suggest that future research should instead focus more on the aspects of the different CMMs and assess methods on a meta-level rather than solely on their performance increases. We further encourage future research to shift its focus toward identifying which tools are best suited for specific CMM-related tasks. This includes not only assessing the cost-benefit trade-offs of available tools but also determining whether a single-tool approach (singular measurement) is sufficient or if a combination of tools (i.e., a multimodal approach) would yield better results.

Finally, while our framework considers important CMMs, future studies should extend it to further metrics that feature prominently in the literature. These include, but are not limited to measures of corporate reputation (Sarstedt et al. 2013) and brand image (Driesner and Romaniuk 2006).

Despite its limitations, our framework offers a first step into leveraging UGC. Researchers and practitioners can draw on the various tools described in our article to put the framework into practice.

Funding The authors did not receive support from any organization for the submitted work.

Conflict of interest R.V. Kübler, S. Adler, L. Welke, M. Sarstedt and K. Pauwels declare that they have no competing interests.

References

Adler, S., and M. Sarstedt. 2021. Mapping the jungle: A bibliometric analysis of research into construal level theory. *Psychology & Marketing* 38(9):1367–1383.

Adler, S.J., M.K. Schöniger, M. Lichters, and M. Sarstedt. 2024. Forty years of context effect research in marketing: A bibliometric analysis. *Journal of Business Economics* 94(3):437–466.

Albers, S. 2010. PLS and success factor studies in marketing. In *Handbook of partial least squares: Concepts, methods and applications in marketing and related fields* Springer Handbooks of Computational Statistics Series, Vol. II, ed. V. Esposito Vinzi, W.W. Chin, J. Henseler, and H. Wang, 409–425. Berlin: Springer.

Anselmsson, J., and N. Bondesson. 2015. Brand value chain in practise; the relationship between mindset and market performance metrics: A study of the Swedish market for FMCG. *Journal of Retailing and Consumer Services* 25:58–70.

Avis, M., S. Forbes, and S. Ferguson. 2014. The brand personality of rocks. *Marketing Theory* 14(4):451–475.

Babić Rosario, A., K. De Valck, and F. Sotgiu. 2020. Conceptualizing the electronic word-of-mouth process: What we know and need to know about eWOM creation, exposure, and evaluation. *Journal of the Academy of Marketing Science* 48(3):422–448.

Baehre, S., M. O'Dwyer, L. O'Malley, and V.M. Story. 2022. Customer mindset metrics: A systematic evaluation of the net promoter score (NPS) vs. alternative calculation methods. *Journal of Business Research* 149:353–362.

Berger, J., A. Humphreys, S. Ludwig, W.W. Moe, O. Netzer, and D.A. Schweidel. 2020. Uniting the tribes: Using text for marketing insight. *Journal of Marketing* 84(1):1–25.
Bergkvist, L., and M. Eisend. 2021. The dynamic nature of marketing constructs. *Journal of the Academy of Marketing Science* 49(3):521–541.
Bergkvist, L., and T. Langner. 2017. Construct measurement in advertising research. *Journal of Advertising* 46(1):129–140.
Bergkvist, L., and T. Langner. 2019. Construct heterogeneity and proliferation in advertising research. *International Journal of Advertising* 38(8):1286–1302.
Blei, D.M., A.Y. Ng, and M.I. Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3:993–1022.
Boegershausen, J., H. Datta, A. Borah, and A.T. Stephen. 2022. Fields of gold: Scraping web data for marketing insights. *Journal of Marketing* 86(5):1–20.
Bollen, K., and R. Lennox. 1991. Conventional wisdom on measurement: A structural equation perspective. *Psychological Bulletin* 110(2):305–314.
Borsboom, D., G.J. Mellenbergh, and J. van Heerden. 2003. The theoretical status of latent variables. *Psychological Review* 110(2):203–219.
Bruner, G.C. 2023. *Marketing scales handbook: Multi-item measures for consumer insight research*. Vol. 12. Independently published.
Burtch, G., Q. He, Y. Hong, and D. Lee. 2022. How do peer awards motivate creative content? Experimental evidence from Reddit. *Management Science* 68(5):3488–3506.
Büschken, J., and G.M. Allenby. 2016. Sentence-based text analysis for customer reviews. *Marketing Science* 35(6):953–975.
Cain, P.M. 2022. Modelling short-and long-term marketing effects in the consumer purchase journey. *International Journal of Research in Marketing* 39(1):96–116.
Carion, N., F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229. Cham: Springer.
Chakravarty, A., Y. Liu, and T. Mazumdar. 2010. The differential effects of online word-of-mouth and critics' reviews on pre-release movie evaluation. *Journal of Interactive Marketing* 24(3):185–197.
Chevalier, J.A., and D. Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research* 43(3):345–354.
Colicev, A., and A. de Bruyn. 2023. The spillover effects of positive and negative buzz on brand attitudes. *European Journal of Marketing* 57(9):2382–2406.
Colicev, A., A. Malshe, K. Pauwels, and P. O'Connor. 2018. Improving consumer mindset metrics and shareholder value through social media: The different roles of owned and earned media. *Journal of Marketing* 82(1):37–56.
Colley, Russel H. 1961. *Defining advertising goals for measured advertising results*. New York: Association of National Advertisers.
Cronbach, L.J. 1951. Coefficient alpha and the internal structure of tests. *Psychometrika* 16(3):297–334.
Daugherty, T., M.S. Eastin, and L. Bright. 2008. Exploring consumer motivations for creating user-generated content. *Journal of Interactive Advertising* 8(2):16–25.
DeVellis, R.F., and C.T. Thorpe. 2021. *Scale development. Theory and applications*, 5th edn., Thousand Oaks: SAGE.
Diamantopoulos, A., and H.M. Winklhofer. 2001. Index construction with formative indicators: An alternative to scale development. *Journal of Marketing Research* 38(2):269–277.
Diamantopoulos, A., M. Sarstedt, C. Fuchs, P. Wilczynski, and S. Kaiser. 2012. Guidelines for choosing between multi-item and single-item scales for construct measurement: A predictive validity perspective. *Journal of the Academy of Marketing Science* 40(3):434–449.
Do, H.H., P.W. Prasad, A. Maag, and A. Alsadoon. 2019. Deep learning for aspect-based sentiment analysis: a comparative review. *Expert Systems with Applications* 118:272–299.
Dotson, J.P., R.R. Fan, E. McDonnell Feit, J.D. Oldham, and Y.-H. Yeh. 2017. Brand attitudes and search engine queries. *Journal of Interactive Marketing* 37(1):105–116.
Driesner, C., and J. Romaniuk. 2006. Comparing methods of brand image measurement. *International Journal of Market Research* 48(6):681–698.
Duarte. 2024. Amount of data created daily. https://explodingtopics.com/blog/data-generated-per-day.
Dzyabura, D., S. El Kihal, J.R. Hauser, and M. Ibragimov. 2023. Leveraging the power of images in managing product return rates. *Marketing Science* 42(6):1125–1142.
Edwards, J.R. 2011. The fallacy of formative measurement. *Organizational Research Methods* 14(2):370–388.

Edwards, J.R., and R.P. Bagozzi. 2000. On the nature and direction of relationships between constructs and measures. *Psychological Methods* 5(2):155–174.
Farris, Paul W., Neil T. Bendle, Philip E. Pfeifer, and David J. Reibstein. 2010. *Marketing metrics: The definitive guide to measuring marketing performance, mergers & acquisitions: The Dealermaker's journal*. Upper Saddle River: Pearson Education.
Fornell, C., M.D. Johnson, E.W. Anderson, J. Cha, and B.E. Bryant. 1996. The American Customer Satisfaction Index: Nature, purpose, and findings. *Journal of Marketing* 60(4):7–18.
Ghosh, A., and M. Deb. 2022. Universal emotional arc structure in advertisements and their favorability. *European Journal of Marketing* 56(4):1232–1267.
Godes, D., and D. Mayzlin. 2004. Using online conversations to study word-of-mouth communication. *Marketing Science* 23(4):545–560.
Gopinath, S., P.K. Chintagunta, and S. Venkataraman. 2013. Blogs, advertising, and local-market movie box office performance. *Management Science* 59(12):2635–2654.
Grootendorst, M. 2022. *BERTopic: Neural topic modeling with a class-based TF-IDF procedure*. arXiv preprint arXiv:2203.05794.
Guyt, J.Y., H. Datta, and J. Boegershausen. 2024. Unlocking the potential of web data for retailing research. *Journal of Retailing* 100(1):130–147.
Guzsvinecz, T., and J. Szűcs. 2023. Length and sentiment analysis of reviews about top-level video game genres on the steam platform. *Computers in Human Behavior* 149:107955.
de Haan, E., M. Padigar, S. El Kihal, R. Kübler, and J.E. Wieringa. 2024. Unstructured data research in business: Toward a structured approach. *Journal of Business Research* 177:114655.
Hair, J.F., and M. Sarstedt. 2021. Data, measurement, and causal inferences in machine learning: Opportunities and challenges for marketing. *Journal of Marketing Theory & Practice* 29(1):65–77.
Hair, J.F., T. Hult, C.M. Ringle, and M. Sarstedt. 2022. *A primer on partial least squares structural equation modeling (PLS-SEM)*, 3rd edn., Thousand Oaks: SAGE.
Hanssens, D.M., K.H. Pauwels, S. Srinivasan, M. Vanhuele, and G. Yildirim. 2014. Consumer attitude metrics for guiding marketing mix decisions. *Marketing Science* 33(4):534–550.
Hartmann, J., M. Heitmann, C. Schamp, and O. Netzer. 2021. The power of brand selfies. *Journal of Marketing Research* 58(6):1159–1177.
Hartmann, J., M. Heitmann, C. Siebert, and C. Schamp. 2023. More than a feeling: Accuracy and application of sentiment analysis. *International Journal of Research in Marketing* 40(1):75–87.
Hennig-Thurau, T., K.P. Gwinner, G. Walsh, and D.D. Gremler. 2004. Electronic word-of-mouth via consumer-opinion platforms: What motivates consumers to articulate themselves on the internet? *Journal of Interactive Marketing* 18(1):38–52.
Henseler, J., C.M. Ringle, and M. Sarstedt. 2015. A new criterion for assessing discriminant validity in variance-based structural equation modeling. *Journal of the Academy of Marketing Science* 43(1):115–135.
Hewett, K., W. Rand, R.T. Rust, and H.J. Van Heerde. 2016. Brand buzz in the echoverse. *Journal of Marketing* 80(3):1–24.
Holiday, S., J.L. Hayes, H. Park, Y. Lyu, and Y. Zhou. 2023. A multimodal emotion perspective on social media influencer marketing: The effectiveness of influencer emotions, network size, and branding on consumer brand engagement using facial expression and linguistic analysis. *Journal of Interactive Marketing* 58(4):414–439.
Hotz-Behofsits, C., N. Wlömert, and Abou N. Nabout. 2025. Natural Affect DEtection (NADE): Using emojis to infer emotions from text. *Journal of Marketing*.
Houston, M.B. 2002. Assessing the validity of secondary data proxies for marketing constructs. *Journal of Business Research* 57(2):154–161.
Howard, John A., and Jagdish N. Sheth. 1969. *The theory of buyer behavior*. New York.
Hulland, J., H. Baumgartner, and K.M. Smith. 2018. Marketing survey research best practices: Evidence and recommendations from a review of JAMS articles. *Journal of the Academy of Marketing Science* 46(1):92–108.
Humphreys, A., and R.J.H. Wang. 2018. Automated text analysis for consumer research. *Journal of Consumer Research* 44(6):1274–1306.
Hunt, H.K. 1991. Consumer satisfaction, dissatisfaction, and complaining behavior. *Journal of Social Issues* 47(1):107–117.
Hydock, C., Z. Chen, and K. Carlson. 2020. Why unhappy customers are unlikely to share their opinions with brands. *Journal of Marketing* 84(6):95–112.
Jebb, A.T., V. Ng, and L. Tay. 2021. A review of key Likert scale development advances: 1995–2019. *Frontiers in Psychology* 12:637547.

Karniouchina, K., K. Sarangee, C. Theokary, and R. Kübler. 2022. The impact of the COVID-19 pandemic on restaurant resilience: Lessons, generalizations, and ideas for future research. *Service Science* 14(2):121–138.
Khodabandehlou, S., and Zivari M. Rahman. 2017. Comparison of supervised machine learning techniques for customer churn prediction based on analysis of customer behavior. *Journal of Systems and Information Technology* 19(1/2):65–93.
Kovács, B. 2024. The Turing test of online reviews: Can we tell the difference between human-written and GPT-4-written online reviews? *Marketing Letters* 35(4):651–666.
Kozinets, R.V., K. De Valck, A.C. Wojnicki, and S.J. Wilner. 2010. Networked narratives: Understanding word-of-mouth marketing in online communities. *Journal of Marketing* 74(2):71–89.
Kübler, R.V. 2023. Will the revolution devour its children? The impact of generative and interactive AI on operative and strategic marketing. *Decisions Marketing* 112(4):267–288.
Kübler, R.V., and S.H. Seggie. 2024. The impact of Covid-19 on how core and peripheral service satisfaction impacts customer satisfaction. *Journal of Business Research* 182:114784.
Kübler, R., K. Pauwels, G. Yildirim, and T. Fandrich. 2018. App popularity: Where in the world are consumers most sensitive to price and user ratings? *Journal of Marketing* 82(5):20–44.
Kübler, R.V., J.E. Wieringa, and K.H. Pauwels. 2017. Machine learning and big data. In *Advanced Methods for Modeling Markets*, 631–670.
Kübler, R.V., A. Colicev, and K.H. Pauwels. 2020. Social media's impact on the consumer mindset: When to use which sentiment extraction tool? *Journal of Interactive Marketing* 50(1):136–155.
Kübler, R.V., L. Lobschat, L. Welke, and H. van der Meij. 2024. The effect of review images on review helpfulness: A contingency approach. *Journal of Retailing* 100(1):5–23.
Lavidge, Robert J., and Gary A. Steiner. 1961. A model for predictive measurements of advertising effectiveness. *Journal of Marketing* 25(6):59–62.
Li, X., M. Shi, and X.S. Wang. 2019. Video mining: measuring visual information using automatic methods. *International Journal of Research in Marketing* 36(2):216–231.
Lin, Y., D. Yao, and X. Chen. 2021. Happiness begets money: Emotion and engagement in live streaming. *Journal of Marketing Research* 58(3):417–438.
Liu, H., I. Chatterjee, M. Zhou, X.S. Lu, and A. Abusorrah. 2020. Aspect-based sentiment analysis: A survey of deep learning methods. *IEEE Transactions on Computational Social Systems* 7(6):1358–1375.
Liu, X., P.V. Singh, and K. Srinivasan. 2016. A structured analysis of unstructured big data by leveraging cloud computing. *Marketing Science* 35(3):363–388.
MacKenzie, S.B., P.M. Podsakoff, and N.P. Podsakoff. 2011. Construct measurement and validation procedures in MIS and behavioral research: Integrating new and existing techniques. *MIS Quarterly* 35(2):293–334.
Matthe, M., D.M. Ringel, and B. Skiera. 2023. Mapping market structure evolution. *Marketing Science* 42(3):589–613.
Mecredy, P.J., M.J. Wright, P.M. Feetham, and P. Stern. 2022. Empirical generalisations in customer mindset metrics. *Journal of Consumer Behaviour* 21(1):102–120.
Mohammad, S.M., and P.D. Turney. 2013. NRC emotion lexicon. *National Research Council, Canada* 2:234.
Onishi, H., and P. Manchanda. 2012. Marketing activity, blogging and sales. *International Journal of Research in Marketing* 29(3):221–234.
Orti, O., R. Tous, M. Gomez, J. Poveda, L. Cruz, and O. Wust. 2019. Real-time logo detection in brand-related social media images. In *Advances in Computational Intelligence: 15th International Work-Conference on Artificial Neural Networks*. IWANN 2019, Gran Canaria, June 12–14, 2019., 125–136. Springer. Proceedings, Part II 15.
Park, C.S., and V. Srinivasan. 1994. A survey-based method for measuring and understanding brand equity and its extendibility. *Journal of Marketing Research* 31(2):271–288.
Pauwels, K. 2014. *It's not the size of the data—it's how you use it: smarter marketing with analytics and dashboards*. Amacom.
Pauwels, K., and B. van Ewijk. 2020. Enduring attitudes and contextual interest: When and why attitude surveys still matter in the online consumer decision journey. *Journal of Interactive Marketing* 52(1):20–34.
Pauwels, K., and A. Joshi. 2016. Selecting predictive metrics for marketing dashboards—an analytical approach. *Journal of Marketing Behavior* 2(2-3):195–224.
Pennebaker, J.W. 2001. *Linguistic inquiry and word count: LIWC 2001*

Petersen, J.A., V. Kumar, Y. Polo, and F.J. Sese. 2018. Unlocking the power of marketing: Understanding the links between customer mindset metrics, behavior, and profitability. *Journal of the Academy of Marketing Science* 46(5):813–836.

Rigdon, E.E., and M. Sarstedt. 2022. Accounting for uncertainty in the measurement of unobservable marketing phenomena. In *Review of Marketing Research*, Vol. 19, ed. H. Baumgartner, B. Weijters, 53–73. Bingley: Emerald.

Rigdon, E.E., J.-M. Becker, and M. Sarstedt. 2019. Factor indeterminacy as metrological uncertainty: Implications for advancing psychological measurement. *Multivariate Behavioral Research* 54(3):429–443.

Rigdon, E.E., M. Sarstedt, and J.-M. Becker. 2020. Quantify uncertainty in behavioral research. *Nature Human Behaviour* 4:329–331.

Ringel, D.M., and B. Skiera. 2016. Visualizing asymmetric competition among more than 1,000 products using big search data. *Marketing Science* 35(3):511–534.

Roberts, M.E., B.M. Stewart, and D. Tingley. 2019. Stm: An R package for structural topic models. *Journal of Statistical Software* 91:1–40.

Rocklage, M.D., D.D. Rucker, and L.F. Nordgren. 2021. Mass-scale emotionality reveals human behaviour and marketplace success. *Nature Human Behaviour* 5(10):1323–1329.

Rust, R.T., W. Rand, M.-H. Huang, A.T. Stephen, G. Brooks, and T. Chabuk. 2021. Real-time brand reputation tracking using social media. *Journal of Marketing* 85(4):21–43.

Salzberger, T., M. Sarstedt, and A. Diamantopoulos. 2016. Measurement in the social sciences: Where C-OAR-SE delivers and where it does not. *European Journal of Marketing* 50(11):1942–1952.

Sarstedt, M., and N.P. Danks. 2022. Prediction in HRM research—A gap between rhetoric and reality. *Human Resource Management Journal* 32(2):485–513.

Sarstedt, M., P. Wilczynski, and T.C. Melewar. 2013. Measuring reputation in global markets—A comparison of reputation measures' convergent and criterion validities. *Journal of World Business* 48(3):329–339.

Sarstedt, M., A. Diamantopoulos, and T. Salzberger. 2016b. Should we use single items? Better not. *Journal of Business Research* 69(8):3199–3203.

Sarstedt, M., A. Diamantopoulos, T. Salzberger, and P. Baumgartner. 2016c. Selecting single items to measure doubly concrete constructs: A cautionary tale. *Journal of Business Research* 69(8):3159–3167.

Sarstedt, M., J.F. Hair, C.M. Ringle, K.O. Thiele, and S.P. Gudergan. 2016a. Estimation issues with PLS and CBSEM: Where the bias lies! *Journal of Business Research* 69(10):3998–4010.

Sarstedt, M., S.A. Adler, L. Rau, and B. Schmitt. 2024a. Using silicon samples in consumer and marketing research: Challenges, opportunities, and guidelines. *Psychology and Marketing* 41(6):1254–1270.

Sarstedt, M., S.J. Adler, C.M. Ringle, G. Cho, A. Diamantopoulos, H. Hwang, and B.D. Liengaard. 2024b. Same model, same data, but different outcomes: Evaluating the impact of method choice in structural equation modeling. *Journal of Product Innovation Management* 41(6):1100–1117.

Schwenzow, J., J. Hartmann, A. Schikowsky, and M. Heitmann. 2021. Understanding videos at scale: How to extract insights for business research. *Journal of Business Research* 123:367–379.

Sepehri, A., M.S. Mirshafiee, and D.M. Markowitz. 2023. PassivePy: A tool to automatically identify passive voice in big text data. *Journal of Consumer Psychology* 33(4):714–727.

TikTok. 2024. API Documentation. https://developers.tiktok.com/products/research-api/on. Accessed 15 Jan 2025.

Toisoul, A., J. Kossaifi, A. Bulat, G. Tziropoulos, and M. Pantic. 2021. Estimation of continuous valence and arousal levels from faces in naturalistic conditions. *Nature Machine Intelligence* 3:42–50.

Vakratsas, D., and T. Ambler. 1999. How advertising works: what do we really know? *Journal of Marketing* 63(1):26–43.

Venkatesan, R., A. Bleier, W. Reinartz, and N. Ravishanker. 2019. Improving customer profit predictions with customer mindset metrics through multiple overimputation. *Journal of the Academy of Marketing Science* 47(5):771–794.

Voulodimos, A., N. Doulamis, A. Doulamis, and E. Protopapadakis. 2018. Deep learning for computer vision: A brief review. *Computational Intelligence and Neuroscience* 2018(1):7068349.

Wang, Q., S. Xing, X. Zhao, and T. Li. 2018. Research on CTR prediction based on deep learning. *IEEE Access* 7:12779–12789.

Wang, X., S. Lu, X.I. Li, M. Khamitov, and N. Bendle. 2021. Audio mining: The role of vocal tone in persuasion. *Journal of Consumer Research* 48(2):189–211.

Washburn, J.H., and R.E. Plank. 2002. Measuring brand equity: An evaluation of a consumer-based brand equity scale. *Journal of Marketing Theory and Practice* 10(1):46–62.

Wu, Y., E.W.T. Ngai, P. Wu, and C. Wu. 2020. Fake online reviews: Literature review, synthesis, and directions for future research. *Decision Support Systems* 132:113280.

X. 2025. API Introduction. https://docs.x.com/x-api/introduction. Accessed 12 Mar 2025.

Xing, X., F. Shi, J. Huang, Y. Wu, Y. Nan, S. Zhang, Y. Fang, M. Roberts, C.-B. Schönlieb, J. del Ser, and G. Yang. 2025. On the caveats of AI autophagy. *Nature Machine Intelligence* 7:172–180.

Yildirim, G., and R. Kübler. 2023. *Applied marketing analytics using R*. Thousand Oaks: SAGE.

Yoo, C., and N. Donthu. 2001. Developing and validating a multidimensional consumer-based brand equity scale. *Journal of Business Research* 52(1):1–14.

Zhou, M., G.H. Chen, P. Ferreira, and M.D. Smith. 2021. Consumer behavior in the online classroom: Using video analytics and machine learning to understand the consumption of video courseware. *Journal of Marketing Research* 58(6):1079–1100.