



Innovative Applications of O.R.

Mining Twitter lists to extract brand-related associative information for celebrity endorsement



Charalampos Saridakis^{a,*}, Constantine S. Katsikeas^a, Sofia Angelidou^b,
Maria Oikonomidou^c, Polyvios Pratikakis^{c,d}

^a Leeds University Business School, University of Leeds, Leeds LS2 9JT, United Kingdom

^b University of Liverpool Management School, University of Liverpool, Liverpool, L69 7ZH, United Kingdom

^c Institute of Computer Science, FORTH, 100N, Plastira St, Vasilika Vouton, Heraklion, GR70013, Greece

^d Computer Science Department, University of Crete, Heraklion, Greece

ARTICLE INFO

Article history:

Received 22 June 2021

Accepted 2 May 2023

Available online 5 May 2023

Keywords:

OR in marketing

Celebrity endorsement

Twitter lists

Big data

Data mining

ABSTRACT

Twitter lists (i.e., curated collections of Twitter accounts) are user-generated and serve primarily as a tool to group other users. Grouping judgments are grounded in the implicit assumption that co-listed members share common associations. As such, Twitter lists are ideal for directly exploring associative links between brands and/or other entities. This research capitalizes on Twitter list membership data to provide a new metric indicating the similarity of users' list membership profiles. This metric is used as a proxy for perceptions of brand–celebrity (mis)fit (i.e., the degree of congruency or similarity between the celebrity and the brand) in celebrity endorsement situations, where a celebrity's fame or social status is used to promote a brand. To validate the accuracy of the method, we compare the list similarity metric with directly elicited survey data for a test set of 62 celebrities and 64 brands, ranging across eight industry sectors. This research contributes to the extant literature of studies extracting brand-related associative information (i.e., information held in consumers' memory that contains the meaning of a brand) from large volumes of consumer online data. This research also introduces new ways of data mining to operational research literature and provides managers with a new methodology to directly infer perceptions of brand–celebrity (mis)fit.

© 2023 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

1. Introduction

Understanding how consumers organize brand information in their memory has been a topic of enduring interest in the marketing literature. The way information is structured in people's minds influences their perceptions of a brand (Ng & Houston, 2009). A brand is defined as a name, term, design, symbol, or any other feature that identifies one seller's goods or service as distinct from those of other sellers (American Marketing Association, 2022). The interpreted meanings of brands are represented by associative structures that consumers generate to link a brand to specific attributes or features, usage situations, or product spokespersons (John, Loken, Kim & Monga, 2006). Keller (2013) defines brand association as an informational node link to a brand node that

is held in consumers' memory and contains the meaning of a brand. It represents the mental connection a customer makes between a brand and other concepts (e.g., attributes, usage situations, product spokespersons) that collectively form a brand's associative structure (Teichert & Schöntag, 2010). For example, Nike is best known by consumers for its “coolness” and “athletic looking” attributes (Dolnicar & Rossiter, 2008), and Stella Artois is perceived as the “wife beater” beer in Britain because of its alleged connection with aggression and binge drinking (Moore, 2007; Wright, 2012). Indeed, a consumer may hold unfavorable perceptions, for instance, when a brand is linked to unpleasant usage situations or a spokesperson who is incongruent with its image. The well-publicized footage of Taliban and al-Qaeda members using Toyota's pickup trucks in their operations has put the firm's reputation at significant risk (Rigby, 2011). Likewise, in 2006, Frédéric Rouzaud, managing director of the Cristal champagne brand, was asked by *The Economist* about the potential threat that Cristal's “bling” association with rappers may pose to the firm's long-term reputa-

* Corresponding author.

E-mail address: b.saridakis@leeds.ac.uk (C. Saridakis).

tion and profitability (Rigby, 2011). Brand associations can provide important clues for tracking brand reputation threats and/or opportunities and can convey how unique the brand is in the eyes of the consumer (Dillon, Madden, Kirmani, & Mukherjee, 2001; Henderson, Iacobucci & Calder, 1998; Torres & Bijmolt, 2009).

From a theoretical perspective, a significant body of work has been dedicated to unraveling how consumers perceive brand associative structures (e.g., Henderson et al., 1998; John et al., 2006; Torres & Bijmolt, 2009). To elicit brand associations and understand how consumers categorize brands, early studies used traditional means of primary data collection, such as consumer surveys, interviews, and focus groups (e.g., Aaker, 1996; Steenkamp & Van Trijp, 1997). More recently, scholars have begun to infer brand associative structures by utilizing publicly available big data, generated by consumers using online platforms (e.g., Culotta & Cutler, 2016; Malhotra & Bhattacharyya, 2022; Nam, Joshi, & Kannan, 2017). As Table 1 shows, research on extracting brand-related associative information from large volumes of consumer data has exploited online search data (e.g., Ringel & Skiera, 2016), microblog data in the form of followers' (e.g., Culotta & Cutler, 2016; Malhotra & Bhattacharyya, 2022; Peng, Agarwal, Hosanagar, & Iyengar, 2018) or tagging (e.g., Nam & Kannan, 2014; Nam et al., 2017) structure, and text-based data from online discussion forums (e.g., Netzer, Feldman, Goldenberg, & Fresko, 2012).

A common characteristic of previous approaches is that brand associative structures and relevant brand categorizations are inferred indirectly. While brand associative structures are the mental connections a customer makes between a brand and other concepts (e.g., attributes, product spokespersons), categorization pertains to a fundamental cognitive ability to recognize shared features, attributes, or similarities among elements (e.g., objects, events, ideas) and to organize these elements by associating them as more abstract groups (i.e., categories) on the basis of these shared attributes, features, or similarities (Croft & Cruse, 2004). Some authors have inferred brand associations from similarities identified between a brand's Twitter follower structure and the follower structure of another brand (Malhotra & Bhattacharyya, 2022) or of an "exemplar" account that has been assumed to strongly represent a given perceptual attribute (Culotta & Cutler, 2016). For example, a similarity between BMW's followers and Green Peace's followers might signal BMW's "eco-friendliness." Other authors categorize brands and infer brand competitive structures from online consumer data pertaining to product- and price-comparison search patterns (Ringel & Skiera, 2016) or infer brand associative information by applying automatic keyword extraction algorithms to text generated in consumers' online discussions (Netzer et al., 2012). In these studies, associative structures are inferred indirectly, as consumers do not explicitly mention brand associations. A few studies analyze social tags, where consumers can explicitly categorize a large volume of content using descriptive words (e.g., "cool," "style," "luxury") and brand identifiers (e.g., "Apple," "Gucci" "Ford") (Nam et al., 2017; Nam & Kannan, 2014). Although a brand's associative structure and categorization can be directly inferred from consumers' explicit tagging activities, this approach does not allow the direct assessment of alignment in the derived categorization. Alignment in categorization refers to an arrangement in which two or more elements (e.g., brands, objects, ideas) are organized in parallel to each other under the same abstract group (i.e., category). For example, while social tags can indicate how a brand is categorized in the consumer's mind, the researcher needs to take an extra step to identify other brands or entities that are also co-aligned in this categorization.

A tool to directly assess alignment in consumers' explicit brand-related categorizations is currently missing from the literature. Given that alignment in categorization between two or more entities may be indicative of perceived fit, if such a tool ex-

isted, it could provide managers with valuable insights into certain branding domains. An example is the branding domain of celebrity endorsement—a form of advertising campaign that involves a well-known person using his or her fame or social status to help promote a branded product or service (Erdogan, 1999). Empirical evidence confirms that higher fit between a brand and a celebrity prompts positive consumer perceptions about the celebrity endorsement partnership (Albert, Ambroise, & Valette-Florence, 2017), increases consumers' willingness to engage with the company (Duthie, Veríssimo, Keane, & Knight, 2017), and reduces the associated financial and reputational risk involved in celebrity endorsement decisions (Carrillat, d'Astous, & Christianis, 2014; Erdogan, 1999) while increasing sales performance (Zheng & Ni, 2020). In contrast, the absence of fit often generates negative connotations because it signals that the brand-celebrity partnership is financially driven (Kamins & Gupta, 1994). Fit refers to the degree of congruency, similarity, resemblance, relevance, or consistency between the celebrity and the brand (Kamins & Gupta, 1994). A good match between a celebrity's image and a brand's image results in more positive consumer perceptions of the advertisement, the celebrity, and the brand than a poor match (Belanche, Casaló, Flavián, & Ibáñez-Sánchez, 2021; Kamins & Gupta, 1994).

Companies spend billions of dollars to hire celebrities that fit their image to endorse their brands. For example, Andy Murray, the British tennis player, has secured an estimated UK£8 million endorsement deal over the next eight years with the sportswear manufacturer Castore (Ahmed, 2019). Likewise, Beyoncé and PepsiCo have a deal estimated to be worth approximately US\$50 million (Sisario, 2012). Such endorsement partnerships have recently expanded to include social media influencers—a business estimated to reach US\$10 billion in 2020 (Kahr, Leitner, Ruthmair, & Sinnl, 2021; Wakabayashi, 2018). A successful endorsement contract with a celebrity can lead to an increase in sales by 4% almost immediately (Bradic, 2015), and the brand's stock has been shown to rise as soon as the news is made public (Crutchfield, 2010; Kraut, Burke, Riedl, & Resnick, 2010). In contrast, negative publicity around a celebrity may transfer harmful associations to the brand. In 2009, for example, Nike lost approximately US\$12 billion due to negative publicity from the marital infidelity scandal of Tiger Woods, who was the brand ambassador for Nike golf apparel and footwear at the time (Knittel & Stango, 2014).

Despite the important role that associative structures play in celebrity endorsement, extant literature has not identified how big data can be used in this particular branding domain (see Table 1). Specifically, marketing scholars have exploited big data to visualize brand market structure and competitive landscape (Netzer et al., 2012), to capture components of customer-based brand equity (Nam & Kannan, 2014), to visualize asymmetric brand competition in product categories (Ringel & Skiera, 2016), to position brands relative to perceptual attributes (Culotta & Cutler, 2016), to visualize heterogeneous brand perceptions (Nam et al., 2017), to predict brand-related content sharing on social media (Peng et al., 2018), and to identify brand extension opportunities (Malhotra & Bhattacharyya, 2022). However, celebrity endorsement scholarship has not yet exploited big data to extract brand-related associative information and explore associative structures, relying instead on mainstream survey data sources and methodological approaches (see, e.g., Parmar, Ghuman, & Mann, 2020; Tian, Tao, Hong, & Tsai, 2021). This raises the question whether big consumer data in online platforms can be employed for the direct assessment of alignment in consumers' explicit brand-celebrity categorizations and, in turn, be used as a proxy for consumer evaluations of celebrity endorsement (mis)fit.

The current study attempts to fill this gap by capitalizing on a popular Twitter activity, namely, Twitter lists. A Twitter list is a curated collection of Twitter accounts. Twitter users can either

Table 1
Comparison with recent studies analyzing brand-related associative information extracted from large volumes of consumer online data.

Study	Branding domain	Associative information extracted	(Mining) method	Platform	Test set	Data set (crawled)	Tool/metric proposed	Validation method
Netzer et al. (2012)	Brand market structure and competitive landscape	Text-pattern co-occurrences in brand-related consumer online discussions	Text mining online consumer-generated discussions	Sedans Forum on Edmunds.com	30 car brands (169 models), described by 1200 common terms	868,174 messages (comprising nearly 6 million sentences), posted by 76,587 unique consumers	A visualization tool illustrating perceived market structure	Survey data elicited from a random sample of 7623 respondents
Nam and Kannan (2014)	Components of customer-based brand equity	Social tag co-occurrence patterns across brands	Mining the social tagging network	Social bookmarking platforms (e.g., Delicious, Digg)	44 brands linked to 7019 key tags	2000 bookmarks collected, and from a total of 60,377 tags, 7019 key tags	Proxy measures capturing components of customer-based brand equity (e.g., brand familiarity, association favorability)	Tested the ability of the derived social tag proxies to predict the brand's actual stock returns
Ringel and Skiera (2016)	Asymmetric brand competition in large markets	Products that are viewed/considered together by the same consumer	Aggregating consumer consideration sets	Product- and price-comparison website	56 television brands, corresponding to 1124 products	Clickstream data of 105,606 consumers	A modeling and two-dimensional mapping tool, visualizing asymmetric competition in product categories	Tested the ability of product- and price-comparison data to predict actual market shares
Culotta and Cutler (2016)	Brand positioning relative to perceptual attributes	Similarity between a brand's and an "exemplar's" account follower structure	Mining the social followers network	Twitter	239 brands and 3 attributes	A total of 30.6 million brand followers, 14.6 million of which were unique	A followers' similarity score, indicating the relationship between a brand and a predefined perceptual attribute	Survey data elicited from a sample of 500 U.S. participants
Nam et al. (2017)	Heterogeneous brand perceptions	Social tag co-occurrence patterns across brands	Mining the social tagging network	Social bookmarking platforms (e.g., Delicious)	7 brands linked to 6000+ key tags	1869 bookmarks and 8610 tags	Aggregate and disaggregate brand perception map	One-on-one interviews with 23 participants to obtain brand concept maps
Peng et al. (2018)	Brand-related content sharing on social media	Similarity between a sender's and a receiver's follower structure	Mining the social followers network	Twitter and Digg	Tweets authored by 9 brands	397 tweets from 12,565 senders and 869,899 receivers, with an average of 8000 followers each	A followers' similarity measure, indicating receivers that are more likely to share content from particular senders	Additional data collected from Digg, pertaining to sharing activities of 31 ads
Malhotra and Bhattacharyya (2022)	Brand extension opportunities	Similarity between a brand's and another brand's follower structure	Mining the social followers network	Twitter	507 brands' Twitter accounts	Brand accounts with between a few thousand and more than a million followers, and a network consisting of 14,000 edges between brands	A followers' similarity score, indicating cross-category brand-brand connections	Survey data elicited from four samples of 250 U.S. participants each
This study	Celebrity endorsement fit and misfit	Similarity between a brand's and a celebrity's account list membership structure	Mining the social list network	Twitter	64 brands and 62 celebrities	881,941 lists, totaling 63.3 million list membership relations—650,406 lists with celebrities, 286,543 with brands, and 55,009 with both celebrities and brands	A list similarity metric, indicating the level of congruence in a brand-celebrity pair	Survey data elicited from samples of 167 (study 1, wave 1), 135 (study 1, wave 2), and 171 (study 2) U.S. participants

create their own lists of selected accounts or subscribe to lists created by other users.¹ The idea behind Twitter lists is to allow users to manage and guide their exposure in the Twitter environment to only the accounts included in their lists. Twitter allows users to create up to 1000 lists, each with up to 5000 users listed on it. A list can be private (i.e., only accessible to the user who created it) or public (i.e., accessible to anyone). Twitter lists, which until now have not been empirically examined, are user-generated and serve primarily as a tool to group other users. The consistent grouping of specific brands/users and celebrities/users is a manifestation of their categorization co-alignment and, in turn, indicative of perceived brand–celebrity congruence. We propose a new method for directly assessing associative links, based on consumer categorizations explicitly displayed in these lists. Our method exploits Twitter list membership data to provide a metric indicating the similarity of the users' list membership profiles. We use this metric as a proxy for perceptions of brand–celebrity (mis)fit in endorsement situations. To validate the accuracy of our method, we compare the proposed list similarity metric with directly elicited, cross-sector survey data pertaining to consumers' perceptions of brand–celebrity (mis)fit. The comparison suggests that the co-membership of two users (e.g., a celebrity and a brand) in several distinct lists reflects shared associations. In other words, our results indicate that brand–celebrity similarity measures mined from Twitter lists can accurately predict consumer evaluations of celebrity endorsement fit as derived from traditional data sources.

Our study makes two important contributions. First, we offer a new tool for mining brand-related associative information from large volumes of consumer online data by capitalizing on Twitter lists. To our knowledge, the proposed tool is the first of its kind in the literature, as it allows for the direct assessment of alignment in categorization. Our method provides unique insights by determining brand associative structures directly displayed by consumers through list-based co-memberships. Thus, we extend the literature by presenting a new method, with unique informational value, for analyzing a large volume of Twitter list membership data (more than 880,000 lists and more than 63 million list membership relationships). Our study responds to recent calls for novel methodologies that identify new online big data sources (e.g., [Kietzmann Paschen, & Treen, 2018](#); [Tirunillai & Tellis, 2014](#); [Zhan & Tan, 2020](#)) and new ways of data mining to increase the effectiveness of operational research (e.g., [Bertsimas, Delarue, Jaillet, & Martin, 2019](#); [Meisel & Mattfeld, 2010](#); [Olafsson, Li, & Wu, 2008](#)).

Second, we demonstrate the value of big data, generated by consumers using online platforms, in the context of celebrity endorsement. This is an important branding domain that has been overlooked in previous studies on mining brand-related associative information. This oversight is significant, as celebrity endorsement offers an attractive setting for developing and testing mining tools that are primarily concerned with how strongly a brand (e.g., Nike) is directly associated in consumers' minds with a particular celebrity (e.g., LeBron James) rather than how strongly a brand (e.g., Nike) is associated in consumers' minds with a specific perceptual attribute (e.g., “coolness,” “athletic looking”). As

such, we make a methodological contribution by showing how publicly available data can be used to capture these connections and, in turn, track complex brand–celebrity associations that occur in consumers' memory. Such an understanding is critical for the design and implementation of effective celebrity endorsement strategies.

2. Theoretical background and related work

2.1. Celebrity endorsement and associative links

Celebrity endorsement is an inherently complex and difficult strategy (e.g., [Erdogan, Baker, & Tagg, 2001](#); [Miciak & Shanklin, 1994](#)). Predictions about an effective match between a celebrity and a desired brand image are usually made via informed judgments about the consistency between the characteristics of the endorser and the attributes of the brand ([Kommiya Mothilal, Mishra, Nishal, Lalani, & Pal, 2022](#); [Misra & Beatty, 1990](#)). This consistency can be crucial in eliciting positive consumer responses to advertising ([Kamins & Gupta, 1994](#); [Knoll & Matthes, 2017](#)). The issue of fit between endorser and brand is an important topic and has been studied extensively under the rubric of the “match-up hypothesis” or the “congruence model” (see, e.g., [Carrilat & Ilicic, 2019](#)).

While such terms are often used interchangeably, they reflect common views about the key underlying elements behind fit. The basic premise underlying these views is that fit stems from an association that links the endorser and the brand ([Johar & Pham, 1999](#); [Rodgers, 2003](#)). This association results from a degree of perceived similarity between the two ([McDaniel, 1999](#)) in terms of image and/or functional attributes ([Rifon, Choi, Trimble, & Li, 2004](#)). The basic principle on which celebrity endorsement research builds is that the more similar two concepts are, the more likely the two concepts will become integrated in an associative network (e.g., [Hamm, Vaitl, & Lang, 1989](#); [Rozin & Kalat, 1971](#)). This associative link underlies perceptions of brand–endorser similarity and drives predicted endorser effects.

In general, much of the research involving associative links employs attribute matching as a salient mode of categorization ([Dolnicar & Rossiter, 2008](#); [John et al., 2006](#)). The widespread assumption is that attribute matching maps onto similarity and forms the basis for category judgments ([Murphy & Medin, 1985](#)). With respect to celebrity endorsement, a celebrity and a brand exhibit categorization alignment if they share common attributes associated with a category ([Rosch, 1978](#)). In simple terms, categorization alignment involves the grouping of celebrities and brands that exhibit common attributes and the exclusion of those that share nothing in common. Attribute matching, in this context, exposes common attributes and serves as a vehicle for mapping grouping preferences ([Shapiro, Spence, & Gregan-Paxton, 2009](#)). Grouping preferences, therefore, constitute an important indicator of celebrity endorsement congruency judgments. We are not alone in emphasizing the importance of grouping preferences in the context of celebrity endorsement. [Murphy and Medin \(1985\)](#) express similar views, and [Spry, Pappu, and Cornwell \(2011\)](#) contend that a brand–celebrity pair is perceived as congruent when its grouping makes sense to the consumer. This simple view suggests that a tool that directly captures consumers' perceptual grouping evaluations will offer the potential to capture brand–celebrity congruency in novel ways. Categorization alignment could be identified, for example, between brands and celebrities that share common attributes through the use of grouping evaluations rather than attribute matching.

To date, however, the literature on celebrity endorsement contains no such tool. Filling this gap will involve revisiting traditional data sources and methods conventionally applied to celebrity endorsement research. Prior studies have typically used surveys em-

¹ Specifically, Twitter users can click the “More” option, which appears under the “Profile” option, in the left-hand side of their Twitter's account homepage. On the dropdown list that appears, users can click “Lists”. A new page will then appear where users can “Discover new Lists”. In Twitter's lists discovery page, users can find recommended lists to follow or search for additional lists by using a search box. Further, Twitter users can create their own lists by clicking the relevant “New List” icon that appears in the upper right corner of the lists discovery page. A “Create a new List” window will then appear where users are asked to choose a name for their list, and a short description. In the same window, Twitter users can also select if they want to make their list private (only accessible to them) or public (anyone can follow the list).

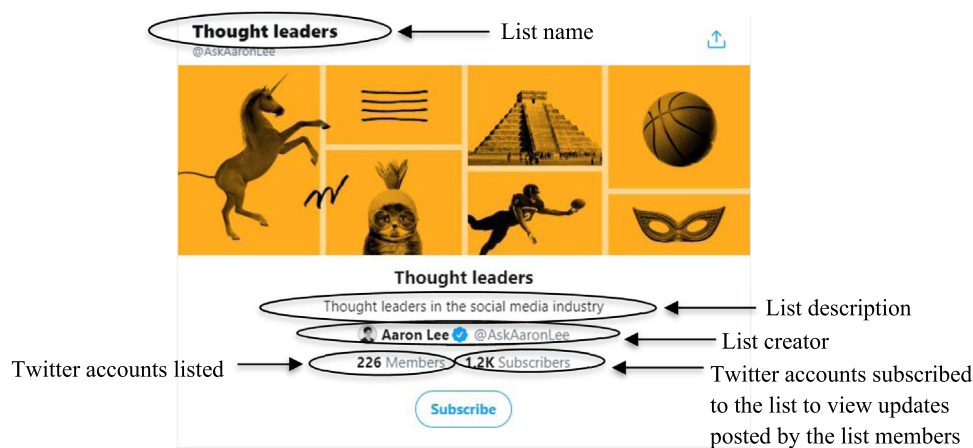


Fig. 1. Example of a Twitter list.

ploying multi-item (e.g., Becker-Olsen & Simmons, 2002; Rifon et al., 2004), nominal, or single-item (e.g., Johar & Pham, 1999) scales. Extant research has paid scant attention to consumer (secondary) data generated in large volumes from online platforms, which can provide rich associative information. We propose that big data generated by users who curate lists on Twitter can provide new insights into consumers' perceptual grouping evaluations. We develop this idea in our discussion of Twitter lists.

2.2. Background for Twitter lists

Twitter is one of the most frequently used social media platforms by celebrities and influencers (Chen, Fan, & Sun, 2019; Pishko, 2019). Twitter introduced lists in 2009, allowing users to group other accounts according to meaningful topics or themes (Greene, O'Callaghan, & Cunningham, 2012; Makrides, Vrontis, & Christofi, 2020). Each Twitter list has a human-readable name, designated by the list creator, which usually indicates the topic that the users in the list are associated with (e.g., "Thought Leaders," "Most Innovative Companies," "Fashion-Beauty"). Twitter lists can also have an optional description, which provides a short definition of the type of users in the list (Kang & Lerman, 2012). Once a list is created, users can add other Twitter accounts (members) without following them directly. In addition, other users can subscribe to the list to view status updates and tweets posted by list members (Zhao & Ram, 2011) (for an example of a Twitter list, see Fig. 1).

Lists serve primarily as a tool to group other user accounts (Golder & Huberman, 2006). From a list creator's viewpoint, the main motivation for creating lists is to unite other users who share some similar underlying attributes (Rakesh, Singh, Vinzamuri, & Reddy, 2014). Users generally have something in common if they are added to the same list (Zhao & Ram, 2011). Regardless of the motivation, the process of list creation involves some type of implicit filtering of the listed users' attributes through existing knowledge structures and schemas. This filtering acts as a basis for categorization (Tomikawa & Dodd, 1980). Thus, the pre-existing knowledge structures and schemas enable the formulation of evaluative inferences about the users' attributes and the subsequent extraction of common associations among the listed users (Maheswaran & Sternthal, 1990).

In this research, we view lists as meaningful manifestations of perceived associative links among listed users. We argue that associative information harvested from Twitter lists can be particularly useful in branding, especially in the context of celebrity endorsement, where managers search for highly congruent celebrities to endorse their brands. These associative links serve as a bridge be-

tween the celebrity and the brand and, in turn, form the basis for transferring additional associations in endorsement situations (e.g., Gregan-Paxton, 2001; Gregan-Paxton & John, 1997). Twitter lists in social networks are increasingly popular and have been used for a variety of purposes in prior empirical literature over the past ten years. For example, they have been used to infer latent characteristics of users (Kim, Moon, & Oh, 2010), discover appropriate topics for a user (Yamaguchi, Amagasa, & Kitagawa, 2011), track well-connected and topic-sensitive followers (Nasirifard & Hayes, 2011), distinguish between elite and ordinary users (Wu, Hofman, Mason & Watts, 2011), measure the degree of homophily between users (Kang & Lerman, 2012), capture emergent semantics (García-Silva, Kang, Lerman, & Corcho, 2012), and uncover relationships between following, membership, and subscription (Velichety & Ram, 2013). More recently, Twitter lists have also been used to provide personalized recommendations (Rakesh et al., 2014), analyze terrorism actions (Kaur, 2016), and identify local communities of users and their common interests (Benabdelkrim, Savinien, & Robardet, 2020). This is the first reported use of Twitter lists in branding to facilitate celebrity endorsement decisions.

2.3. Data-mining methods

Data mining has become an increasingly important area in the operational research literature (e.g., Baesens, Mues, Martens, & Vanthienen, 2009; Meisel & Mattfeld, 2010; Olafsson, Li, & Wu, 2008). Data-mining methods can be broadly classified into two categories: text mining and social network mining (Oikonomidou, Pratikakis, Saridakis, & Angelidou, 2019). Text-mining methods refer to the process of extracting interesting and nontrivial patterns of knowledge from unstructured text documents (Gaikwad, Chaugule, & Patil, 2014; Sumathy & Chidambaram, 2013; Symitsi, Stamolampros, Daskalakis, & Korfiatis, 2021). Text-based analysis has attracted considerable attention in the marketing literature, especially for extracting information from user-generated content such as consumer reviews. Various issues have been investigated, including the impact of consumer reviews on sales (Berger, Sorensen, & Rasmussen, 2010), the relative importance of reviews compared with own experience in consumers' learning process about products (Zhao, Yang, Narayan, & Zhao, 2013), the most significant characteristics being discussed in customer reviews (Shama & Dhage, 2018), the change in conversion rates as a result of changes in affective content and linguistic style of online reviews (Ludwig et al., 2013), the prediction of product sales based on review content and sentiment (Godes & Mayzlin, 2004), the elicitation of product attributes and consumer preferences (Lee & Bradlow, 2011), and

the conversion of online discussions to market structure insights (Netzer et al., 2012).

However, when inferring consumers' perceptions from social media data, text-mining methods limit researchers to harvesting information about topics that are commonly discussed in social media platforms. In contrast, by focusing on platform-based social network structures, rather than on the text generated by users, we can capture unique information that often remains unharvested when relying solely on text data (Culotta & Cutler, 2016). Because our aim in this study is to elicit perceptions of congruence in brand–celebrity pairs, a topic which many users may be less likely to discuss explicitly in online text, social network structures can be particularly useful. Indeed, researchers have already begun to use online social network data for a variety of marketing purposes. For example, social network mining has been used to predict consumer behavior (Goel & Goldstein, 2014), to understand information diffusion (Goel, Watts, & Goldstein, 2012), to measure social media interaction (Yu & Hu, 2020), and to elicit brand-related associative information (e.g., Culotta & Cutler, 2016; Malhotra & Bhattacharyya, 2022; Nam et al., 2017), which is our focus in this research.

Given the constraints associated with using text-mining methods in social media platforms such as Twitter, we explore a new, unexploited source of information, namely, the social network of curated lists. This is the first study to use Twitter lists' social networks as a measure of brand–celebrity congruence perceptions. In addition to capturing important information from a large, broad-based sample of users, investigating Twitter lists' social networks produces a measure that directly assesses alignment in categorization, a particularly important issue when managers are evaluating various celebrities as potential endorsers of their brands. The proposed list-based approach has several advantages over existing methods for eliciting brand-related associative information from consumer data, especially in the domain of celebrity endorsement. Table 2 provides a summary comparison of the list-based approach with existing methods.

Primary data-based approaches are typically designed to collect highly tailored information and reveal unconscious aspects of consumers' perceptions. However, they tend to be more costly, difficult to administer, and sensitive to sample size, and they often depend on the researcher's interpretation. Text mining employs automatic keyword extraction algorithms (Netzer et al., 2012) that can be sensitive to a researcher's assumptions. Because perceptions of congruence in brand–celebrity pairs are a topic that most consumers are less likely to discuss directly in their online text, relevant associative structures and derived categorizations can be inferred only indirectly. In social network–mining methods, such as the followers-based approach (Culotta & Cutler, 2016), the associative structure is indirectly inferred from the similarity between brand followers and “exemplar” account followers. Thus, just like text mining, the followers-based approach does not allow for the direct inference of categorization. In contrast, both the tag-based (e.g., Nam et al., 2017) and the list-based approaches that we propose allow for a direct inference of categorization. Nonetheless, the list-based approach is more appropriate in the context of celebrity endorsement, as it also permits a direct assessment of alignment in this categorization.

3. Proposed mining method from Twitter account data

3.1. Twitter corpus

Drawing on existing Twitter list memberships, we can construct a list–user graph, where an edge between a list and a user node indicates that the list contains the specified user (i.e., brand or celebrity). As an example, Fig. 2 shows a simple graph that depicts three lists. The users “Brand B” and “Celebrity C” are co-listed

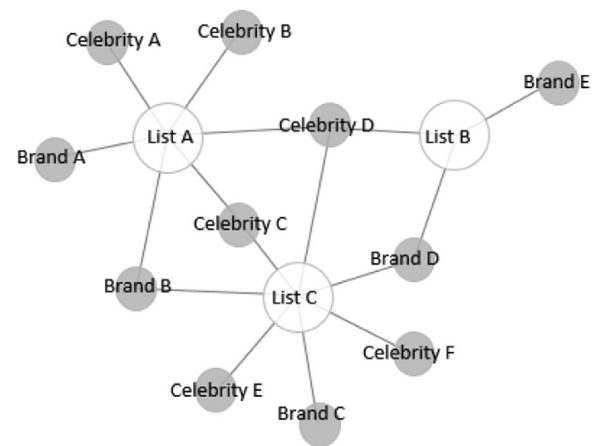


Fig. 2. Illustrative user-list graph containing 11 users (brand and celebrity accounts) grouped into three Twitter lists.

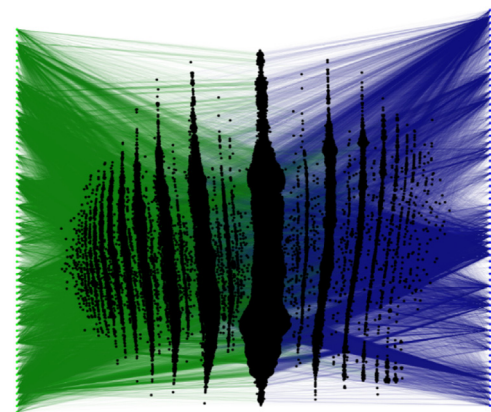


Fig. 3. Illustration of the data set's similarity cloud.

twice because both users are members of lists A and C. Replicating this co-listing across the wider Twitter network may be indicative of an affinity between the pair of users. Using a list–user matrix representation, we can compute a measure that indicates the similarity of the users' list membership profiles. Users who are more frequently co-listed are deemed to be more similar.

To assess the usefulness of information harvested from Twitter list networks for identifying highly congruent celebrities as brand endorsers, we first constructed a Twitter data set based on a test set of targeted Twitter user accounts. The targeted accounts are the official accounts of our chosen test set of brands and celebrities. In total, we targeted 64 brand Twitter accounts and 62 celebrity Twitter accounts (we detail the process for the construction of our test set in the “Validation Method” section). For each targeted account, we identified all Twitter lists to which they have been added by crawling the Twitter API (application programming interface) between December 2019 and February 2020. For each identified list, we also crawled Twitter's list object to collect information on the creator of the list, the title and description of the list, and the number of memberships and subscriptions to the list. This approach yielded a data set of 881,941 lists (i.e., 826,932 lists containing either brands or celebrities, and 55,009 lists containing both brands and celebrities) and 126 targeted users (i.e., 64 brands, 62 celebrities), totaling 63,270,277 list membership relationships.²

Fig. 3 shows the derived similarity cloud for our data set, with brands (highlighted in blue), celebrities (highlighted in green), and

² The full Twitter API code is available from the authors upon request.

Table 2
Comparison of alternative methods for eliciting brand-related associative information from consumer data.

Method	Data collection effort	Data volume	Richness of elicited information	Cost	Relevance to the branding domain of this study
Primary data (e.g., in-depth personal interviews, focus groups, structured questionnaires)	High (-) (design of instruments, training of data collectors, recruitment of study participants, etc.)	Low (-) (data from a sample of few, representative subjects)	High (+) (tailored information, unconscious aspects can be revealed)	High (-) (qualitative/quantitative analysis expert required, incentive for respondent participation needed, etc.)	High (+) Cross-sector survey data also collected for validation purposes.
Text mining (Netzer et al., 2012)	Moderate (-/+) (automated text classification by a text-mining tool, based on a predefined keyword list)	High (+) (large-scale data)	Moderate (-/+) (constrained by algorithmic interpretation)	Moderate (-/+) (multiple stages of text-mining processes)	Low (-) Perceptions of congruence in brand–celebrity pairs is a topic most users are less likely to directly discuss in online text. Associative structure can be inferred only indirectly.
Social network mining: followers (Culotta & Cutler, 2016)	Low (+) (directly stated by consumers/online users, easily available as secondary data)	High (+) (large-scale data)	Moderate (-/+) (depending on the problem, it may be limited and sensitive to customers' biases and potential social influences)	Low (+) (publicly available and readily accessible)	Low (-) Associative structure is indirectly inferred from the similarity between brand followers and "exemplar" account followers. Therefore, followers' structure does not allow for the direct inference of categorization.
Social network mining: social tags (Nam et al., 2017; Nam & Kannan, 2014)	Low (+) (directly stated by consumers/online users, easily available as secondary data)	High (+) (large-scale data)	Moderate (-/+) (depending on the problem, it may be limited and sensitive to customers' biases and potential social influences)	Low (+) (publicly available and readily accessible)	Moderate (-/+) Tags are appropriate for the direct inference of categorization, but they do not allow for the direct assessment of alignment in this categorization.
Social network mining: lists (this study)	Low (+) (directly stated by consumers/online users, easily available as secondary data)	High (+) (large-scale data)	Moderate (-/+) (depending on the problem, it may be limited and sensitive to customers' biases and potential social influences)	Low (+) (publicly available and readily accessible)	High (+) Lists represent the most appropriate crowdsourced method for the direct assessment of alignment in categorization.

lists that contain both celebrities and brands (highlighted in black) depicted as graph nodes. Each list is linked to at least one brand and at least one celebrity in the graph. Depending on the ratio of brands/celebrities in a list, the list is placed closer to brands or closer to celebrities, while “superstar” brand or celebrity nodes (i.e., popular brand or celebrity accounts included in a relatively larger number of lists) are depicted with darker blue or darker green shaded edges, respectively. A quantization of lists into sets of equal brand-to-celebrity ratios emerges.

To visualize more clearly the similarity networks derived from our proposed method, we present an illustrative example. Drawing on existing Twitter list memberships, we generate a list-user graph depicting the resulting associative network of three brand users and three celebrity users from our test set (Fig. 4).

Brands (highlighted in blue), celebrities (highlighted in green), and lists (highlighted in black) are depicted as graph nodes in Fig. 4. An edge between a list and a brand or celebrity node indicates that the list contains the specified user. Each list in the graph contains at least one brand, at least one celebrity, or both. “Superstar” brand or celebrity nodes are depicted with darker blue or darker green shaded edges, respectively, and are also surrounded by larger black-shaded areas (i.e., lists in which these users are not co-listed with any other users). Brand-user accounts appear to be more popular than celebrity-user accounts, as they appear in more lists.

The list-user graph in Fig. 4 suggests, for example, that the user “Porsche” (automobile brand) is co-listed with the user “Team Messi” (football player) in more lists compared with the number of lists that “Porsche” is grouped with the user “garthbrooks”




Fig. 4. List-user graph depicting three brand users and three celebrity users from our test set.

(songwriter). This may be indicative of a higher affinity between “Porsche” and “Team Messi” than between “Porsche” and “garthbrooks.” Similar conclusions can be drawn for “American Express” (financial services company) and “Jeff Immelt” (business executive) as opposed to “American Express” and “Team Messi.”

Relying on the identified similarities of the users' list membership profiles, we can compute a relevant measure indicating the level of congruence between pairs of users, namely, brands and celebrities. In the following section, we turn to this computation.

3.2. List similarity metric

We propose a list similarity metric, which extracts perceived correspondences from list membership data. As we have explained,

developing this metric is based on the idea that, when Twitter users create lists by selecting specific accounts to be included as members, list creators implicitly identify some underlying common associations shared by all listed members. In other words, these list memberships are not random, but instead reflect an implicit filtering of characteristics and attributes according to the creators' knowledge structures and schemas. These knowledge structures and schemas help list creators formulate categorization judgments based on some common underlying properties shared among all accounts that the creator adds to the list. On several occasions, Twitter users explicitly group sets of accounts into topical or other categories on the basis of common interests, expertise, and other characteristics. For example, a Twitter user may decide to group family members in a Twitter list. It is logical to believe that such a grouping judgment is grounded in the implicit assumption that all co-listed family members share some common associations (e.g., they are members of the same family; have strong emotional bonds with each other; and/or share common characteristics, experiences, interests, etc.). This renders Twitter lists ideal for directly exploring associative links between members that are co-listed. We aim to harvest this underlying associative information of perceived similarity between targeted accounts (i.e., celebrities and brands) by considering the number of lists in which the targeted accounts are placed together. We assert that when an independent Twitter user, who voluntarily spends time to create lists, has listed two targeted accounts (e.g., a brand and a celebrity) together, the accounts must have something in common (Kim et al., 2010). For example, Messi (football player) and Porsche (automobile brand) are co-listed in a Twitter list named “Prospects”—a list that provides its followers with updates on the latest current affairs posted by the listed members on a variety of topics, including the latest sports-related news. Thus, our central hypothesis is as follows: The higher the number of independent lists in which two targeted accounts—in our case, a brand's and a celebrity's Twitter account—are placed together, the stronger is the similarity (i.e., perceived fit) between these accounts, and thus, the greater is the potential for effective endorsement. Moreover, the co-membership of two accounts in several lists means that several users have independently identified them as having some similar underlying property (Benabdelkrim et al., 2020).

However, we acknowledge that the co-membership count of the targeted accounts is often skewed, following the “superstar effect” (Rosen, 1983) or the “winner-take-all” phenomenon (Frank & Cook, 2010). Specifically, it is particularly common to the Twitter graph that the distribution of list membership follows a power law, meaning that a few extremely popular accounts are included in a large number of lists. To account for this, we normalize our proposed similarity metric using a common and empirically successful similarity function, the Jaccard index—also known as the Jaccard similarity coefficient (Jaccard, 1908). This index, defined as the size of the intersection divided by the size of the union of the sample sets, is a widely used statistic for understanding the similarities between sample sets. The mathematical representation of our Jaccard-normalized list similarity metric (L_{sim}) between brand b and celebrity c is specified as follows:

$$L_{sim}(b, c) = \frac{|L(b) \cap L(c)|}{|L(b) \cup L(c)|},$$

where $L(b)$ and $L(c)$ are all the lists that include brand b and celebrity c , respectively.

To compute the Jaccard-normalized list similarity metric, we analyzed 63,270,277 list membership relationships in our data set, counting the number of lists that include both brand b and celebrity c , divided by the total list membership of brand b and celebrity c , for all selected pairs (b, c). These lists include those

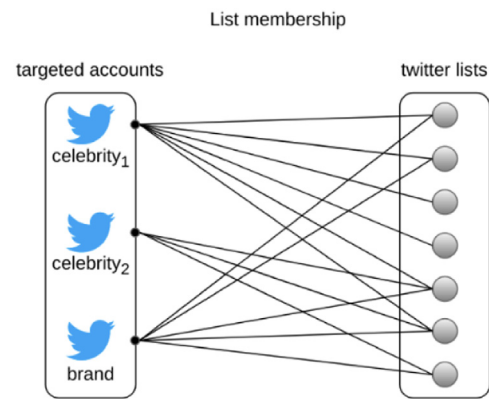


Fig. 5. Illustration of the Jaccard-normalized list similarity.

created and curated by users in no way linked to b or c , not just lists by followers or friends of b and c .

The Jaccard index is often used in social network analysis because it is normalized over the set of total list membership, thus enabling a comparison of accounts included in different numbers of lists. For example, as Fig. 5 shows, the Twitter account of celebrity₁ (C_1) is placed in more lists than that of celebrity₂ (C_2) and thus has more lists in common with any brand (B) simply because of its larger sample (superstar effect). The Jaccard index adjusts for this bias by normalizing over the total lists membership of the celebrity and brand, resulting in $L_{sim}(C_1, B) < L_{sim}(C_2, B)$.

For the illustrative example presented in Fig. 5, the Jaccard-normalized list similarity metric for the brand (B) with celebrity₁ (C_1) would yield $L_{sim}(C_1, B) = 4/7 = 0.57$, whereas with celebrity₂ (C_2), it would be equal to $L_{sim}(C_2, B) = 3/5 = 0.6$. This suggests that $L_{sim}(C_1, B) < L_{sim}(C_2, B)$, even though $L(C_1) \cap L(B) > L(C_2) \cap L(B)$.

Finally, our Jaccard-normalized list similarity metric also accounts for the presence of potential negativity bias in curated lists—an issue that arises when people and brands are evaluated unfavorably (Alves, Koch, & Unkelbach, 2017; Rozin & Royzman, 2001). The negativity bias is a cognitive bias suggesting that, even when of equal intensity, things of a more negative nature (e.g., unpleasant thoughts, emotions, or information) have a greater effect on one's psychological state and decisions than neutral or positive things (Rozin & Royzman, 2001). An extension of this bias is that individuals are typically more likely to continue to interact with people if they have a positive impression of them, while they keep people they dislike at a distance (e.g., Denrell, 2005; Fazio, Eiser, & Shook, 2004). In the same vein, it is reasonable to expect that Twitter users tend to keep people (e.g., celebrities) and brands they dislike at a distance, and therefore they may have limited information or lack well-established schemas about these celebrities and brands. For example, there are brand-celebrity pairs that are co-listed in lists named “mortal-enemies” or “things-I-hate.” There is a possibility that any grouping activity in such negatively polarized lists may be driven by isomorphism in the negative feelings toward the brand and celebrity that are co-listed rather than characteristic- or attribute-related similarities identified by the creators of the relevant lists. This possibility suggests that celebrities and brands co-listed in such negatively polarized lists may be less appropriate endorsement matches.

To identify and control for targeted user pairs that appear in lists with a negative valence, we followed several steps. Because every Twitter list has a name and a description attribute, we first performed sentiment analysis on these to discover the polarity of each association. Sentiment analysis is implemented with respect to a body of text to understand the sentiment expressed in it. Typically, this sentiment is quantified with a positive or negative

value, or a polarity (Li, Gupta, Zhang, & Flor, 2020). The overall sentiment is inferred from the sign of the polarity score. Based on this score, the body of text is classified as positive, neutral, or negative according to the overall sentiment expressed by its author (Molina-González, Martínez-Cámara, Martín-Valdivia, & Perea-Ortega, 2013). We performed sentiment analysis using lexicon-based models. In lexicon-based sentiment analysis, words in texts are labelled as positive, negative, or neutral with the help of a valence dictionary. Once each word in the text is labelled, an overall sentiment score is then derived by counting the numbers of positive and negative words and then combining these values mathematically (a combining function, usually the sum or average, is taken to make the final prediction regarding the overall sentiment).

We repeated the analysis for every list in which the targeted brands and celebrities co-appear. Traditionally, sentiment analysis computes the strength of the sentiment along with its polarity, most often resulting in sentiment vectors that include positive and negative scores for a piece of text (Ilk & Fan, 2022; Mai & Le, 2021; Sul, Dennis, & Yuan, 2017). To prevent lists with stronger wording from disproportionately affecting our final sentiment scores, we focus on the polarity of the sentiment alone. The sole aim of this filtering mechanism is to identify whether a list has a positive or negative connotation—not how strong this connotation is according to the view of a single list curator. Second, we computed the number of positive and negative lists for each pair. If the number of positive lists is equal to or higher than the number of negative lists, then the overall valence for this pair is positive (or negative otherwise).

The two attributes of each list that may contain sentiment—namely, the title and the description—often include the use of language that differs in style and form. The title usually consists of a few keywords (on average, 2 to 4), which are often connected with a hyphen and do not form a full sentence. On the contrary, descriptions are more often free-form text, which include full sentences and punctuation and may be the length of a small paragraph. Due to the different nature of the language used in each attribute, we employ two separate lexicons for the sentiment analysis. Specifically, we use AFINN—a dictionary of words that rates connotation—to compute the sentiment score for the title attribute. We chose AFINN for the title attribute because it is a lexicon developed primarily to analyze sentiment in very small Twitter texts by assigning each word a sentiment score (Nielsen, 2011). We use Vader—a lexicon and rule-based sentiment analysis tool (Darko & Liang, 2023; Hutto & Gilbert, 2014)—to conduct sentiment analysis on the description attribute. We chose Vader for the description attribute because it is a lexicon specifically attuned to sentiments expressed in free-form text on social media. In addition, Vader can gauge overall syntactical sentiment—making it more appropriate for the description attribute, which includes full sentences—while AFINN can assess types of words used—making it more appropriate for the title attribute, which consists of a few keywords (Chappelka, Oh, Scott, & Walker-Holmes, 2017).

After conducting sentiment analysis on both attributes, we end up with two polarity “votes” for each Twitter list. We then count and aggregate these votes separately. Assume, for example, that a celebrity (C) is placed on four lists, a brand (B) is placed on seven lists, and the celebrity (C) and brand (B) are placed together on three lists. Their Jaccard-normalized list similarity score is as follows: $L_{sim}(C, B) = 3/8 = 0.375$. After calculating the similarity score, we perform sentiment analysis on the three common lists to obtain two separate polarity scores for each list. In our example, we assume that we get the following results (+1 = positive connotation, 0 = neutral sentiment, -1 = signifies negative connotation):

- List₁: title score = +1, description score = 0,
- List₂: title score = 0, description score = -1,

- List₃: title score = +1, description score = +1.

Subsequently, we aggregate these scores to get the overall sign for each pair, which we then multiply by L_{sim} . If a list has an empty description field, we count 0 on the description score. In our example, the final similarity score is given by the following function:

$$L_{sim}(C, B) \times \text{sign}(\text{List}_1\text{scores} + \text{List}_2\text{scores} + \text{List}_3\text{scores}) \\ = +L_{sim}(C, B) = +0.375.$$

We can now infer that the list similarity score between the celebrity (C) and the brand (B) is positive, with a 0.375 value of similarity strength.

We use the overall sign of the sentiment score—rather than the sentiment score itself—as the weight, for two main reasons. First, to account for the presence of potential negativity bias in curated lists, our intention in this context is to identify and control for all the targeted user pairs that appear in lists with a negative valence, regardless of how strong this valence may be. Second, the sign of the sentiment score is more trustworthy than the score itself. Specifically, we use the sign of the sentiment to equalize how much each list affects the final similarity score. Had we not done so, a list title or description using strong or explicit language would affect the similarity score much more than a list title or description using polite language, even though both lists may express the same sentiment. To avoid this bias, especially because opinions on Twitter are frequently very polarized, we consider each list to amount to “a vote of opinion” of equal weight.

4. Validation method

The next step is to validate the extent to which the proposed automated list similarity metric matches actual consumer perceptions of fit or congruence in brand–celebrity pairs. To achieve this, we compare automated list similarity metric scores with directly elicited survey similarity responses for brand–celebrity pairs included in our test set.

4.1. Brand and celebrity selection for the test set

Our fixed targeted users are the official Twitter accounts (see, e.g., Paul, Khattar, Chopra, Kumaraguru, & Gupta, 2019) of a test set that includes brands and celebrities. We chose the shortlisted brands and celebrities in the test set as follows: First, using official 2017 *Forbes*’ ranking lists, we identified the 100 “World’s Most Valuable Brands.” Similarly, celebrity targets were identified by merging two *Forbes*’ personality-related ranking lists—namely, the 2017 *Forbes*’ list of the 100 “World’s Most Powerful People” and the 100 “World’s Highest-Paid Celebrities.” By merging these two personality lists, we can generalize our approach across different personality types that can be used as brand endorsers (e.g., athletes, entertainers, dignitaries). Second, we manually matched the identified targeted brands and personalities with their official Twitter account. We dropped brands or personalities maintaining no or more than one official Twitter accounts from the lists. Third, to test the generalizability of our mining method across sectors, we grouped all remaining targeted brands under eight industry sectors (i.e., automobiles and parts, financial services, food and beverage, media and telecommunications, personal and household goods, retail, technology, and industrial goods and services), using the Industry Classification Benchmark—a taxonomy launched by Dow Jones and Financial Times Stock Exchange in 2005. Classification of brands into these sectors relied on information provided on their official websites. Furthermore, to account for the presence of negative valence on social media, which could artificially inflate the results of our proposed similarity measure, the actual test set

Table 3
Illustrative targeted brands by sector.

Sector	Illustrative brands	N	Illustrative occupations of personalities used in the consideration sets	N (validation study 1)
Automobiles and parts	Toyota, Ford, Audi, Mercedes Benz	8	Football player, entrepreneur, singer and songwriter, basketball player	40
Financial services	American express, Mastercard, Citi, JP Morgan	8	Media personality, journalist, basketball player, actor and film producer	40
Food and beverage	Starbucks, Coca-Cola, Corona, Nestlé	8	Radio personality, golfer, technology entrepreneur, talk show host	40
Media and telecommunications	Facebook, Netflix, Disney, Fox TV	8	Singer and songwriter, stand-up comedian, tennis player, journalist	40
Personal and household goods	Colgate, Gillette, NIVEA, L'Oreal Paris	8	Racing driver, actor, entrepreneur, baseball pitcher	40
Retail	H&M, Walmart, ZARA, CVS Pharmacy	8	Producer, writer and filmmaker, actor, basketball player	40
Technology	Philips, Huawei, Microsoft, Intel	8	Stand-up comedian, baseball player, court judge, electronic DJ and producer	40
Industrial goods and services	Caterpillar, Boeing, General Electric, FedEx	8	Economist, media personality, singer, football player	40

consisted of target accounts with no prior involvement in major scandals. Finally, we excluded any brands carrying a full or partial celebrity name. This selection process resulted in a test set of 64 brands and 62 personalities.

4.2. Validation study 1

In study 1, the 64 targeted brands were grouped into eight balanced and sector-specific subsets (i.e., eight brands per sector). A randomly generated consideration set was allocated to each brand, which included five personalities from the original pool of 62 personalities in the test set. This random allocation resulted in 40 potential brand–celebrity pairs per industry sector (i.e., 8 brands per sector \times 5 randomly allocated celebrities in its relevant consideration set) and 320 potential brand–celebrity pairs across all eight sectors. In constructing the 64 consideration sets for the targeted brands, we ensured that none of the randomly allocated personalities in a given choice set had engaged in any type of endorsement relationship in the past with the allocated brand. Broadly, we accounted for the presence of “inorganic” similarity in each of the 320 potential brand–celebrity pairs (e.g., carefully coordinated moves, induced by the firm and/or the celebrity to engage in similar content or strategy in their Twitter accounts) that could potentially bias the results of our proposed similarity measure. Specifically, we carried out relevant online searches via the three most popular search engines (i.e., Google, Bing, Yahoo) to ensure that the pairs do not co-appear or relate in any way to material published online.

Table 3 provides illustrative examples of targeted brands per sector. Within each sector, we included different types of personalities in the consideration sets. We also provide the occupation of some illustrative targeted personalities.

Next, using our test set of 64 brands, we directly elicited survey responses to determine how well consumers associate each brand with each of the five randomly allocated personalities in the respective consideration set. We administered the survey through Amazon Mechanical Turk (AMT), which has been shown to be a reliable source for social science data collection (e.g., [Buhmester, Kwang, & Gosling, 2011](#)). Indeed, recent work on the usefulness of AMT in academic research suggests that AMT samples are more representative of the general population than participants from other subject pools (e.g., [Hulland & Miller, 2018](#); [Matherly, 2019](#); [Smith, Roster, Golden, & Albaum, 2016](#)). To assess our sample's representativeness, we collected key demographic data on participant characteristics. The median income of participants is \$62,389, which maps to the U.S. median household income of \$61,937 ([Guzman, 2019](#)). Similar to the Twitter population, participants are

relatively young, with a median age of 36.5 years, which is representative of the U.S. population median age of 38.2 years ([U.S. Census Bureau, 2019](#)). Consistent with the gender distribution of the Twitter population (i.e., 65% of Twitter audiences are male and 35% are female) ([Statista, 2020](#)), our sample has a higher representation of males (64.3%) than females (35.7%). To further increase the validity of responses, we required participants to be based in the United States and to have a successful track record on AMT (i.e., completion of at least 100 prior assignments with an acceptance rate of at least 95%). After grouping the test set's 64 brands into eight subsets by sector, we asked AMT participants to rate each subset.

In study 1, we aimed to validate our list similarity scores by approaching data collection as a real managerial problem. Specifically, for each brand, participants were asked to choose the personality who would provide the best fit and those who would provide a bad fit, if they knew that the relevant brand manager had short-listed the given five personalities in the consideration set as potential endorsers. Participants could select a separate column if they did not recognize a brand or a celebrity. Although we collected relevant choice data for all 64 brands in our test set, to make the tasks more manageable for respondents and increase the validity of responses, we collected data in two waves. In each wave, participants were randomly allocated to identify the best and bad fits for half of the brands in the test set (i.e., in each wave, each participant evaluated 32 brands, 4 per sector). In both waves, we randomized brand and personality order in the consideration sets for each participant. We also included attention filters to ensure valid responses by asking participants to specify the brand they had evaluated earlier, to select a particular response for a given line, and/or to appropriately indicate that they did not recognize a nonsense word inserted in place of a real brand. In addition, survey respondents had to be active Twitter users. We discarded responses for those who did not pass these checks, resulting in a total sample of 167 (wave 1) and 135 (wave 2) respondents.

4.3. Study 1 results

To evaluate the effectiveness of the proposed automated list similarity metric in capturing fit across sectors, we report the percentage of agreement between our list similarity metric and participants' responses. Table 4 presents these results. Across all sectors, the list similarity metric correctly predicted respondents' celebrity choices of best fit in 76.56% of tasks—81.25% and 71.88% in wave 1 and wave 2, respectively. Respondents' celebrity choices for bad fit were correctly predicted in 94.53% of tasks. Correct predictions for best fit range from 50% (food and beverage sector) to 100% (fi-

Table 4
Validation study 1 (waves 1 and 2): percentage of correct celebrity choice predictions.

Sector	% Correctly identified celebrity with best fit	% Correctly identified celebrity with bad fit	Chi-square (p-value)	N (number of tasks)
Automobiles and parts	50.00	90.63	7.32 (0.02)	40
Financial services	100.00	100.00	40.00 (0.00)	40
Food and beverage	50.00	87.50	5.63 (0.04)	40
Media and telecommunications	87.50	96.88	28.48 (0.00)	40
Personal and household goods	87.50	96.88	28.48 (0.00)	40
Retail	62.50	90.63	11.29 (0.00)	40
Technology	87.50	96.88	28.48 (0.00)	40
Industrial goods and services	87.50	96.88	28.48 (0.00)	40
Average	76.56	94.53	163.67 (0.00)	320

Table 5
Validation study 2: percentage of correct brand–celebrity pairing predictions.

Sector	% Correctly identified pairs (good or bad pairs)	Chi-square (p-value)	N (number of tasks × individuals)
Automobiles and parts	71.64	128.09 (0.00)	684
Financial services	76.32	189.47 (0.00)	684
Food and beverage	65.79	68.21 (0.00)	684
Media and telecommunications	80.70	257.90 (0.00)	684
Personal and household goods	78.07	215.58 (0.00)	684
Retail	88.30	401.43 (0.00)	684
Technology	83.33	304.00 (0.00)	684
Industrial goods and services	76.61	193.71 (0.00)	684
Average	77.60	1666.74 (0.00)	5472

financial services sector), while correct predictions for bad fit range from 87.50% (food and beverage sector) to 100% (financial services sector). The proposed metric performed well, correctly predicting survey responses across all sectors and brands approximately eight out of ten times for celebrities with the best fit and nine out of ten times for celebrities with bad fit. Correct choice predictions were relatively lower within the automobile and parts and the food and beverage sectors (best fit was correctly predicted in 50% of the tasks, while bad fit was correctly predicted in 90.63% and 87.50% of tasks, respectively). Chi-square tests suggest that all identified differences between predicted and survey-elicited responses were statistically significant ($p < 0.05$).

4.4. Validation study 2

In study 2, as an additional analysis, we randomly picked only a few brands (i.e., eight brands in total, or one brand per sector) and determined the relevant list similarity scores between them and all 62 celebrities in our original test set. Contrary to study 1, in study 2 we focus on a few brands in the test set, as pairing all 64 brands with all 62 celebrities would have resulted in an unmanageable number of pairs. We ranked all celebrities in our test set on our list similarity metric and found the best two and the worst two overall celebrity pairings for each of the eight brands ($2 \times 8 + 2 \times 8 = 32$ pairs in total). As in study 1, in study 2 we validated our list similarity scores by approaching data collection as a real managerial problem. We informed respondents that these pairs had been short-listed by relevant brand managers and asked respondents to evaluate them and identify good and bad pairings. Similar to study 1, we ensured that none of the pairs had engaged in any type of endorsement relationship in the past, by accounting for the presence of “inorganic” similarity.

We administered the survey via AMT, and each of the recruited participants evaluated all 32 “good” and “bad” pairs. Participants could select a separate column if they did not recognize a brand or a celebrity in a given pair. Participants also had to be based in the United States and have a successful track record on AMT. We randomized the order of appearance of the pairs for each participant. As in study 1, we included the same attention filters to en-

sure valid responses. Likewise, survey respondents had to be active Twitter users. We eliminated participants who did not meet these criteria, resulting in a total sample of 171 respondents. The final set of questions focused on respondent demographics, including gender, marital status, age, education, occupation, and annual personal income.

4.5. Study 2 results

To evaluate the effectiveness of the proposed automated list similarity metric in capturing survey-elicited perceived brand–celebrity fit, we report the level of agreement between survey responses and list similarity metric scores (Table 5). Across all sectors, respondents identified pairs with good and bad fit consistently with predictions based on our list similarity scores 77.60% of the time. Correct predictions range from 65.79% (food and beverage sector) to 83.33% (technology sector). Chi-square tests suggest that all identified differences between predicted and survey-elicited responses were statistically significant ($p < 0.01$).

The proposed metric performed well in both celebrity choice (Study 1) and brand–celebrity pairing (Study 2) tasks, correctly predicting survey responses across all sectors and brands eight out of ten times. Because automated brand–celebrity similarity perception estimation is a novel contribution to the marketing literature and the stream of studies eliciting brand-related associative information from big data, there is no clear external benchmark against which to directly compare the performance of our approach. For an indirect comparison, we consider the studies of Netzer et al. (2012) and Culotta and Cutler (2016), who also use survey data to validate their approaches. Netzer et al. (2012) present an automated data methodology for estimating car brand co-association sets (i.e., car brands that consumers cluster together in purchase consideration sets), using text analysis of online user forum posts. In their validation, they obtained correlations between their estimates and survey results ranging from 0.43 to 0.55. Likewise, Culotta and Cutler (2016), using a metric based on Twitter followers to capture perceptual brand attribute ratings, obtained correlations ranging from 0.62 to 0.82. Given that these studies examine brand market structure and perceptual attribute ratings as opposed

Table 6
Validation study 1 (waves 1 and 2): percentage of correct celebrity choice predictions for higher-order relationships.

Sector	% Correctly identified celebrity with best fit	% Correctly identified celebrity with bad fit	Chi-square (p-value)	N (number of tasks)
Automobiles and parts	11.11	25.00	0.57 (0.45)	40
Financial services	50.00	25.00	5.63 (0.01)	40
Food and beverage	0	0	2.50 (0.11)	40
Media and telecommunications	25.00	12.50	0.15 (0.69)	40
Personal and household goods	11.11	12.50	0.35 (0.55)	40
Retail	0	12.50	2.50 (0.11)	40
Technology	25.00	37.50	0.15 (0.69)	40
Industrial goods and services	50.00	25.00	5.62 (0.01)	40
Average	21.54	18.46	0.12 (0.73)	320

Table 7
Validation study 2: percentage of correct brand–celebrity pairing predictions for higher-order relationships.

Sector	% Correctly identified pairs (good or bad pairs)	Chi-square (p-value)	N (number of tasks × individuals)
Automobiles and parts	28.36	0.00 (1.00)	684
Financial services	47.37	0.00 (1.00)	684
Food and beverage	34.21	0.00 (1.00)	684
Media and telecommunications	51.46	4.00 (0.05)	684
Personal and household goods	21.93	0.00 (1.00)	684
Retail	47.66	4.00 (0.05)	684
Technology	50.87	0.00 (1.00)	684
Industrial goods and services	47.95	4.00 (0.05)	684
Average	41.22	4.50 (0.03)	5472

to our focus on perceptual brand–celebrity similarity, any comparisons need to be made with caution. This said, the evidence cited here compares well with that of prior research efforts. Finally, we performed two additional robustness checks to verify that our predictions are indicative of brand–celebrity similarity associative information and not driven by popularity or demographic differences (for details, see Online Appendix A).

5. Supplemental analyses

Our proposed list similarity metric relies on first-order relationships, in the sense that brand–celebrity pairs that are not directly co-listed in any lists receive a score of “0” (meaning no similarity or fit). An example is the brand A–celebrity E pair in Fig. 2. In social networks, however, researchers may also be interested in examining higher-order relationships. Such relationships suggest that two objects are similar if they are linked with objects that are themselves similar (Jeh & Widom, 2002). For example, although there are no edges between any single list and the brand A and celebrity E nodes in Fig. 2, there are nodes—specifically, brand B, celebrity C, and celebrity D—that serve as bridges between list A and list C that contain brand A and celebrity E, respectively. Higher-order relationships may suggest that brand A and celebrity E nodes are similar, even if they are not first-order neighbors, as they are linked to the interconnected lists A and C, respectively.

To examine whether the consideration of such “higher-order neighbors” improves the ability to predict brand–celebrity (mis)fit perceptions, we performed a supplemental analysis.³ Specifically, we computed the SimRank metric, which is a general similarity measure that is applicable to domains that are naturally modelled as graphs, with nodes representing objects (e.g., brands, celebrities, lists) and edges representing relationships among them (Jeh & Widom, 2002). Moreover, it is an edge metric in that it estimates a single similarity score for each pair of nodes (e.g., brand–celebrity pair), and therefore, it is directly comparable to our proposed list similarity metric. The SimRank is a higher-order metric, in that it

is affected by the structure of the full graph, not just the relationship of the immediate (direct) neighbors of the brand and celebrity nodes (Chen, Fan, & Sun, 2021; Jeh & Widom, 2002).

The basic SimRank equation is given as follows: For a node v in a graph, we denote the set of in-neighbors of v and out-neighbors of v as $I(v)$ and $O(v)$, respectively. Individual in-neighbors are denoted as $I_i(v)$, for $1 \leq i \leq |I(v)|$, and individual out-neighbors are denoted as $O_i(v)$, for $1 \leq i \leq |O(v)|$. We denote the similarity between brand b and celebrity c as $s(b, c) \in [0, 1]$. We can then write a recursive equation for $s(b, c)$. If $b = c$, then $s(b, c)$ is defined to be 1. Otherwise,

$$s(b, c) = \frac{C}{|I(b)||I(c)|} \sum_{i=1}^{|I(b)|} \sum_{j=1}^{|I(c)|} s(I_i(b), I_j(c)),$$

where C is a constant between 0 and 1. A minor technicality here is that either b or c may not have any in-neighbors. Because there is no way to infer any similarity between b and c in this case, similarity is set to $s(b, c) = 0$, so the summation in the equation is defined to be 0 when $I(b) = \emptyset$ or $I(c) = \emptyset$.

For each brand–celebrity pair in our test set, we compute a SimRank score, which captures the structural similarity of nodes within the network recursively (Jeh & Widom, 2002). For a direct comparison against our proposed list similarity metric, we performed the same validations by examining how well the SimRank metric predicts survey responses for our test set of 62 celebrities and 64 brands. Validation results suggest that our proposed list similarity metric outperforms the SimRank metric in both celebrity choice (Study 1) and brand–celebrity pairing (Study 2) tasks. On average, the SimRank metric correctly predicted survey responses only two out of ten times across sectors and brands in Study 1 (see Table 6) and four out of ten times across sectors and brands in Study 2 (see Table 7).

From a theoretical perspective, focusing on first-order relationships alone to infer perceptions of brand–celebrity (mis)fit seems more relevant and informative in the given context domain than extending to higher-order relationships. First-order neighbors, which appear in several lists together, explicitly share some common underlying attributes that unite them and form the basis for

³ We thank an anonymous reviewer for bringing this point to our attention.

their direct co-categorization. In contrast, this is not necessarily the case with higher-order neighbors. For example, the user “Estée Lauder” is unlikely to have a high affinity with the user “Usain Bolt,” even if there is a node (e.g., “Maria Sharapova”) that serves as a bridge between the lists “beautiful world” and “famous athletes” that contain Estée Lauder and Usain Bolt, respectively. It is logical to believe that the higher the distance between two neighbors, the weaker are the inferences the researcher can make about their common underlying attributes and thus their level of affinity. Finally, an additional weakness of the SimRank and other similar higher-order metrics (e.g., PageRank [see Page, Brin, Motwani, & Winograd, 1999; Scholz, Pfeiffer, & Rothlauf, 2017]) relative to our proposed metric is that they are computationally more expensive and require the full network for computation. In contrast, our list similarity metric can be computed on demand, using partial data targeting a specific brand–celebrity pair. In our context, for example, we computed the SimRank metric only on the subset of the co-membership network consisting of brand and celebrity nodes due to its high computational cost.

Last, to examine whether sentiment analysis improves the ability of our proposed metric to predict brand–celebrity (mis)fit perceptions, we performed an additional supplemental analysis. Specifically, for each brand–celebrity pair in our test set, we compute a list similarity metric score that does not incorporate the sentiment analysis component. Subsequently, we performed the same validations by examining how well the list similarity metric without the sentiment analysis component predicts survey responses for our test set of 62 celebrities and 64 brands. Validation results suggest that our proposed list similarity metric with sentiment analysis outperforms the metric that does not incorporate the sentiment analysis, in both celebrity choice (Study 1) and brand–celebrity pairing (Study 2) tasks. On average, the metric that does not incorporate the sentiment analysis component correctly predicted survey responses seven out of ten times across sectors and brands, in both studies—eight out of ten times was the relevant score based on the proposed metric with the sentiment analysis component. Although our proposed metric performs well in both cases, the sentimental analysis component slightly improves its performance. From a theoretical perspective, individuals keep objects they dislike at a distance (e.g., Denrell, 2005; Fazio, Eiser, & Shook, 2004) and, hence, they may possess limited information about such objects (e.g., celebrities and brands) that are co-listed in negatively polarized lists.

6. Discussion and implications

6.1. Contribution to knowledge

Twitter lists, a common activity on Twitter, provide valuable insights into whether consumers using the platform perceive chosen consideration sets of celebrities to have a good or bad fit with different brands. The power of Twitter lists in our context resides in the interpretation that users generate with respect to shared associations between a celebrity and a brand. Our results, which are robust to alternative research design settings, indicate that the co-membership of two Twitter users (e.g., celebrity and brand) in several distinct lists is a good reflection of these shared associations.

Our research contributes to the stream of studies focusing on the elicitation of brand-related associative information from big data by proposing a new method. We propose a Twitter list-based approach that can be employed to directly infer brand associative structures and categorization. Our work is the first to provide a tool that quantifies the degree of alignment in consumers' explicit categorizations and to investigate its accuracy in capturing perceptions of fit. This easy-to-implement and fully automated metric can be more effective in certain branding domains than existing tools,

as it goes beyond brand-related attributional matches to directly assess how tightly two or more entities (e.g., brands, celebrities) are co-aligned in consumers' associative structures. While prior approaches have paved the way for developing new data-mining methods that can assist several marketing tasks, they are more complicated and time consuming than our approach because they require prior tracking and specification of common dimensions or attributes (e.g., Culotta & Cutler, 2016; Nam et al., 2017). Thus, with existing methods, associative structures and alignment in categorization can be inferred only indirectly—a premise that makes them less effective in capturing perceptions of (mis)fit.

Given the unique informational value of our proposed proxy to accurately capture the strength of association between selected entities, we show how it serves as a new, efficient, and precise method to track perceptions of fit or congruence in the context of celebrity endorsement decisions. To our knowledge, this is the first attempt to address this issue, not only in the area of celebrity endorsement but also in the broader marketing literature. Extant research has focused on problem domains in which category judgments can be tracked effectively using brand-related attributional matching (e.g., Nam & Kannan, 2014; Ringel & Skiera, 2016). Using celebrity endorsement as a focal branding domain, the current study provides a more complete picture of how big data can be used to infer brand-related associative information. We provide evidence that big data can also be used for observing brand–celebrity associative links and offer new insights into grouping preferences concerning celebrity endorsement congruency judgments.

We validate the notion that grouping preferences can accurately predict perceptions of brand–celebrity fit, confirming prior work suggesting that consumers perceive brand–celebrity pairs as congruent when their grouping makes sense to them (e.g., Murphy & Medin, 1985; Spry et al., 2011). Unlike attribute matching, which has been used for the analysis of similarity judgments in other problem domains, grouping preferences can deal with multiple attributes simultaneously (Ben-Av & Sagi, 1995). This is particularly important, given that any set of two entities may share a seemingly infinite number of common attributes. Thus, the actual verdict on whether two entities are perceived to be more or less similar depends heavily on the particular weights of salient importance that each consumer allocates to each individual attribute (Murphy & Medin, 1985). For example, Kim Kardashian might be perceived as being more (less) similar to IBM than Christopher Langan, if the attribute “rich” is more (less) important than the attribute “smart” in the consumer's mind. In essence, the accurate prediction of similarity judgments on the basis of attribute matching requires prior knowledge of the importance weights that consumers allocate to each attribute. In contrast, grouping preferences seem to be mediated by attribute selection in the sense that grouping information implies that two or more entities are co-categorized on the basis of preselected and shared salient attributes (e.g., Huang, 2015; Levinthal & Franconeri, 2011). In other words, a particular consumer will be able to group Twitter users together only after devoting attention to different attributes that are particularly important in their background knowledge. This opens new avenues for research on brand association through the lens of similarity judgments as a combination of multiple attribute cues.

Our measure provides a way to directly assess similarity, regardless of the complexity and multiplicity of attribute cues that come into play during consumers' similarity judgments. At the most fundamental level, our approach inverts the process of capturing perceived similarities, from identifying common attributes that potentially justify subsequent similarity judgments to directly evaluating the final similarity judgments *per se*. This shift from attribute matching to a more generative attribute selection view of similarity expands the scope of brand associative structures in the context of celebrity endorsement. Such a view calls for attention

to meaningful similarity judgment outcomes, which emerge from a constellation of shared attributes that are salient in the consumer's memory and that, in turn, have significant implications for the transferring of additional associations in endorsement situations (e.g., [Gregan-Paxton, 2001](#); [Gregan-Paxton & John, 1997](#)).

6.2. Managerial implications

Our list similarity metric provides brand managers with a new and easily applied approach to accurately assess perceived brand–celebrity fit from publicly available online data. Although such data frequently contain valuable information about consumers' brand-related perceptions ([Buzeta, De Pelsmacker, & Dens, 2020](#)), marketers often find the task of mining and analysis challenging and complex. Our method provides a roadmap for extracting information from big data in online platforms—in this case, from Twitter lists. Marketing practitioners can use this method to automatically evaluate existing or potential celebrity endorsement decisions for a large set of personalities of interest. Such evaluations can be useful for a brand to map alternative celebrity endorsement options that have the potential to initiate positive associations in endorsement situations and to disregard options that might have detrimental reputation effects.

Because celebrity endorsements are expensive for the brand, our tool could be particularly useful to managers because it helps identify large sets of congruent celebrities for targeting and thus can help managers select more attractive and affordable options. While similar insights can be obtained using survey-based elicitation methods, our method provides similarity scores for an unprecedented number of brand–celebrity pairs, which gives it several advantages over traditional methods. This capability provides managers with the unique opportunity to evaluate multiple brand–celebrity pairs at a low cost quickly and effectively. Our approach is attractive due to its accuracy, simplicity, and computational tractability.

Twitter lists curated by individual consumers may also provide marketing researchers with a unique opportunity to observe consumers' heterogeneous perceptions of congruence between brands and celebrities. Because not all groups of users who generate lists are equally important for brands, an analysis of consumers' heterogeneous congruence perceptions using clustering techniques can help marketing managers discover distinct associations and implement tailored celebrity matches to access targeted consumer segments. An investigation into the dynamic nature of list curation could reveal how brand–celebrity congruence perceptions change over time following significant incidents (e.g., a new product launch). As content (e.g., product reviews, newspaper articles, blogs, commentaries) emerges for a brand in response to an incident, a potential shift in the type of celebrities with which the brand is co-listed may reveal interesting insights. As more content emerges on the incident, its aftermath and firm actions will influence a brand's list membership profile over time. Information about the temporal effect of incidents on list membership profiles can be useful for managers to understand how customer perceptions change over time and how to effectively adapt celebrity endorsement strategies as a result of these changes. In Online Appendix B, we provide a nine-step guide for managers on how to integrate our approach into the process of choosing celebrities as endorsers for their brands.

6.3. Limitations and future research directions

This research and our proposed metric should be considered in the context of certain limitations. First, our sample may not be representative of the entire population of current and prospective customers, due in part to the existence of platform-specific

dynamics and sample bias (e.g., [Schweidel & Moe, 2014](#)) common across all studies that use social media data ([Ruths & Pfeffer, 2014](#)). This may explain the lack of full overlap of our measure with the survey-based consumer metrics (see [Pauwels & Van Ewijk, 2013](#)). Because Twitter list data rely on user input, users' self-selection plays a role in their participation decision on relevant platforms and their choice of brands or celebrities to list. For example, those who list a particular brand and/or celebrity might be more knowledgeable about them and more engaged with news and content related to them. Furthermore, the characteristics of list curators can depend on the characteristics of the particular platform. Twitter users tend to be younger, and the majority are male. To tackle this self-selection bias and obtain more representative data, it is critical for marketing managers to understand whether characteristics of their aggregate customer base are in line with the characteristics of list curators. Further, Twitter allows its users to keep their lists private. This implies that our data set does not include private lists, and thus the examined relationships are based on membership data crawled from public lists. Although the majority of Twitter lists are public, we understand that focusing on public lists alone may bias our findings. Future work might collect data from multiple platforms, thus controlling for platform-specific dynamics that pose threats to the generalizability of findings.

Second, our method does not account for social signaling effects, which may artificially inflate our similarity metric in some sectors. Specifically, in sectors with a high potential for status signaling, users may be more motivated to use Twitter lists to engage with relevant brands than in sectors with a low potential for status signaling. This is a possible explanation for some of our findings. For example, in two high-status-signaling sectors (e.g., automobiles and parts, food and beverages), our similarity choice predictions were lower than predictions in other industries. Future studies should control for these signaling effects when interpreting the co-membership of users across different sectors.

Third, we did not ask survey participants whether they have listed or followed on Twitter the brands or celebrities included in the test set. We expect that the percentage of positive responses for any given brand in this context would be too small to conduct meaningful analysis. However, researchers might explore more direct connections between survey respondents and Twitter activity. Furthermore, although our method is straightforward and easy to implement, future work might supplement our technique with text-mining and data reduction methods (e.g., [Puranam et al., 2017](#); [Tirunillai & Tellis, 2014](#)) to infer specific common attributes reflected in Twitter lists. For example, because each Twitter list has a name indicating the topic or theme that the co-listed users pertain to, this reflects the implicit rationale for their grouping (e.g., "Thought Leaders," "Most Innovative Companies"). The examination of the rationales for grouping in Twitter lists is a fruitful avenue for further research, as this could provide brand managers with rich qualitative insights into the representative attributes that describe their brands and the celebrity selected each time.

Fourth, a natural extension of our study would be to apply this methodology to other substantive marketing tasks. For example, empirical evidence from several studies confirms that higher fit is related to more positive brand–product extension evaluations (e.g., [Aaker & Keller, 1990](#); [Boush & Loken, 1991](#)) and higher rates of sponsor–sponsee partnership success (e.g., [Johar, Pham, & Wakefield, 2006](#); [Olson & Thjomøe, 2009](#)). Instead of estimating brand–celebrity similarities, researchers could compute brand–brand, brand–product, and sponsor–sponsee similarities and use them to identify competitive market structures (e.g., [Henderson et al., 1998](#); [Netzer et al., 2012](#); [Urban, Johnson, & Hauser, 1984](#)), brand extension opportunities ([Aaker & Keller, 1990](#); [Park, Milberg, & Lawson, 1991](#)), and areas for potential

sponsorships (e.g., Johar et al., 2006; Olson & Thjomøe, 2009), respectively.

Finally, the approach we advance in this work can serve as a foundation for future research developments capable of exploiting new sources of big data from social media platforms (e.g., likes, shares, mentions, impressions, URL clicks, comments, hashtag usage, keyword usage) in marketing. For example, by examining “likes” (i.e., approvals of someone else’s content) and “shares” (i.e., reposts of someone else’s content), researchers may be able to assess the level of consumers’ engagement with a brand and/or identify influencers viewed as experts within their niche. A comparison of a brand’s social media followers with those of competitive brands, along with an examination of the type of content these competitor followers engage with, could reveal opportunities and transformative trends in a given market.

CRedit authorship contribution statement

Charalampos Saridakis: Supervision, Methodology, Formal analysis, Conceptualization, Writing – original draft, Writing – review & editing. **Constantine S. Katsikeas:** Writing – review & editing. **Sofia Angelidou:** Methodology, Formal analysis, Conceptualization, Writing – original draft, Writing – review & editing. **Maria Oikonomidou:** Methodology, Formal analysis, Writing – original draft. **Polyvios Pratikakis:** Methodology, Formal analysis, Writing – original draft, Writing – review & editing.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.ejor.2023.05.004.

References

- Aaker, D. A. (1996). Measuring brand equity across products and markets. *California Management Review*, 38(3), 102–121.
- Aaker, D. A., & Keller, K. L. (1990). Consumer evaluations of brand extensions. *Journal of Marketing*, 54(1), 27–41.
- Ahmed, M. (2019). Tennis star Andy Murray secures £8 m sponsorship deal. *Financial Times*. <https://www.ft.com/content/db2513f6-400b-11e9-b896-fe36ec32aece>. (retrieved November 23, 2020).
- Albert, N., Ambrose, L., & Valette-Florence, P. (2017). Consumer, brand, celebrity: Which congruency produces effective celebrity endorsements? *Journal of Business Research*, 81, 96–106.
- Alves, H., Koch, A., & Unkelbach, C. (2017). Why good is more alike than bad: Processing implications. *Trends in Cognitive Sciences*, 21(2), 69–79.
- American Marketing Association (2022), “American Marketing Association dictionary—Official definition of *Brand*,” <https://www.ama.org/topics/branding/> (retrieved April 10, 2022).
- Baesens, B., Mues, C., Martens, D., & Vanthienen, J. (2009). 50 years of data mining and OR: Upcoming trends and challenges. *Journal of the Operational Research Society*, 60(sup1), S16–S23.
- Becker-Olsen, K., & Simmons, C. J. (2002). When do social sponsorships enhance or dilute equity? Fit, message source, and the persistence of effects. *Advances in Consumer Research*, 29(1), 287–289.
- Belanche, D., Casaló, L. V., Flavián, M., & Ibáñez-Sánchez, S. (2021). Understanding influencer marketing: The role of congruence between influencers, products and consumers. *Journal of Business Research*, 132, 186–195.
- Benabdelkrim, M., Savinien, J., & Robardet, C. (2020). Finding interest groups from Twitter lists. In *Proceedings of the 35th annual ACM symposium on applied computing* (pp. 1885–1887). <https://doi.org/10.1145/3341105.3374077>.
- Ben-Av, M. B., & Sagi, D. (1995). Perceptual grouping by similarity and proximity: Experimental results can be predicted by intensity autocorrelations. *Vision Research*, 35(6), 853–866.
- Berger, J., Sorensen, A. T., & Rasmussen, S. J. (2010). Positive effects of negative publicity: When negative reviews increase sales. *Marketing Science*, 29(5), 815–827.
- Bertsimas, D., Delarue, A., Jaillet, P., & Martin, S. (2019). Travel time estimation in the age of big data. *Operations Research*, 67(2), 498–515.
- Boush, D. M., & Loken, B. (1991). A process-tracing study of brand extension evaluation. *Journal of Marketing Research*, 28(1), 16–28.
- Bradic, L. (2015). Celebrity endorsements on social media are driving sales and winning over fans. *Social Media Week*. <https://socialmediaweek.org/blog/2015/09/brands-using-celebrity-endorsements/>. (retrieved July 8, 2019).
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon’s Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1), 3–5.
- Buzeta, C., De Pelsmacker, P., & Dens, N. (2020). Motivations to use different social media types and their impact on consumers’ online brand-related activities (COBRAs). *Journal of Interactive Marketing*, 52, 79–98.
- Carrillat, F. A., d’Astous, A., & Christianis, H. (2014). Guilty by association: The perils of celebrity endorsement for endorsed brands and their direct competitors. *Psychology & Marketing*, 31(11), 1024–1039.
- Carrillat, F. A., & Illicic, J. (2019). The celebrity capital life cycle: A framework for future research directions on celebrity endorsement. *Journal of Advertising*, 48(1), 61–71.
- Chappelka, M., Oh, J., Scott, D., & Walker-Holmes, M. (2017). Food for thought: Analyzing public opinion on the supplemental nutrition assistance program. arXiv preprint arXiv:1710.02443.
- Chen, Z. Y., Fan, Z. P., & Sun, M. (2019). Individual-level social influence identification in social media: A learning-simulation coordinated method. *European Journal of Operational Research*, 273(3), 1005–1015.
- Chen, Z. Y., Fan, Z. P., & Sun, M. (2021). Tensorial graph learning for link prediction in generalized heterogeneous networks. *European Journal of Operational Research*, 290(1), 219–234.
- Croft, W., & Cruse, D. A. (2004). *Cognitive linguistics*. Cambridge, UK: Cambridge University Press.
- Crutchfield, D. (2010). Celebrity endorsements still push product: Why, in the era of social media, the rewards continue to outweigh the risks. *AdAge*. <https://adage.com/article/cmo-strategy/marketing-celebrity-endorsements-push-product/146023>. (retrieved August 8, 2019).
- Culotta, A., & Cutler, J. (2016). Mining brand perceptions from Twitter social networks. *Marketing Science*, 35(3), 343–362.
- Darko, A. P., & Liang, D. (2023). A heterogeneous opinion-driven decision-support model for tourists’ selection with different travel needs in online reviews. *Journal of the Operational Research Society*, 74(1), 272–289.
- Denrell, J. (2005). Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review*, 112(4), 951–978.
- Dillon, W. R., Madden, T. J., Kirmani, A., & Mukherjee, S. (2001). Understanding what’s in a brand rating: A model for assessing brand and attribute effects and their relationship to brand equity. *Journal of Marketing Research*, 38(4), 415–429.
- Dolnicar, S., & Rossiter, J. R. (2008). The low stability of brand-attribute associations is partly due to market research methodology. *International Journal of Research in Marketing*, 25(2), 104–108.
- Duthie, E., Verissimo, D., Keane, A., & Knight, A. T. (2017). The effectiveness of celebrities in conservation marketing. *PLoS One*, 12(7), Article e0180027.
- Erdogan, B. Z. (1999). Celebrity endorsement: A literature review. *Journal of Marketing Management*, 15(4), 291–314.
- Erdogan, B. Z., Baker, M. J., & Tagg, S. (2001). Selecting celebrity endorsers: The practitioner’s perspective. *Journal of Advertising Research*, 41(3), 39–48.
- Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology*, 87(3), 293–311.
- Frank, R. H., & Cook, P. J. (2010). *The winner-take-all society: Why the few at the top get so much more than the rest of us*. New York: Penguin.
- Gaikwad, S. V., Chaugule, A., & Patil, P. (2014). Text mining methods and techniques. *International Journal of Computer Applications*, 85(17), 42–45.
- García-Silva, A., Kang, J. H., Lerman, K., & Corcho, O. (2012). Characterising emergent semantics in Twitter lists. In *Extended semantic web conference* (pp. 530–544). Springer.
- Godes, D., & Mayzlin, D. (2004). Using online conversations to study word-of-mouth communication. *Marketing Science*, 23(4), 545–560.
- Goel, S., & Goldstein, D. G. (2014). Predicting individual behavior with social networks. *Marketing Science*, 33(1), 82–93.
- Goel, S., Watts, D. J., & Goldstein, D. G. (2012). The structure of online diffusion networks. In *Proceedings of the 13th ACM conference on electronic commerce* (pp. 623–638).
- Golder, S. A., & Huberman, B. A. (2006). Usage patterns of collaborative tagging systems. *Journal of Information Science*, 32(2), 198–208.
- Greene, D., O’Callaghan, D., & Cunningham, P. (2012). Identifying topical Twitter communities via user list aggregation. arXiv preprint arXiv:1207.0017.
- Gregan-Paxton, J. (2001). The role of abstract and specific knowledge in the formation of product judgments: An analogical learning perspective. *Journal of Consumer Psychology*, 11(3), 141–158.
- Gregan-Paxton, J., & John, D. R. (1997). Consumer learning by analogy: A model of internal knowledge transfer. *Journal of Consumer Research*, 24(3), 266–284.
- Guzman, G. (2019). *US median household income up in 2018 from 2017*. U.S. Census Bureau <https://www.census.gov/content/dam/Census/library/publications/2019/acs/acsbr18-01.pdf> (retrieved September 10, 2020).
- Hamm, A. O., Vaitl, D., & Lang, P. J. (1989). Fear conditioning, meaning, and belongingness: A selective association analysis. *Journal of Abnormal Psychology*, 98(4), 395–406.
- Henderson, G. R., Iacobucci, D., & Calder, B. J. (1998). Brand diagnostics: Mapping branding effects using consumer associative networks. *European Journal of Operational Research*, 111(2), 306–327.
- Huang, L. (2015). Grouping by similarity is mediated by feature selection: Evidence from the failure of cue combination. *Psychonomic Bulletin & Review*, 22(5), 1364–1369.
- Hulland, J., & Miller, J. (2018). Keep on Turkin? *Journal of the Academy of Marketing Science*, 46, 789–794.

- Hutto, C. J., & Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the eighth international aaii conference on weblogs and social media*. AAAI Press.
- Ilk, N., & Fan, S. (2022). Combining textual cues with social clues: Utilizing social features to improve sentiment analysis in social media. *Decision Sciences*, 53(2), 320–347.
- Jaccard, P. (1908). Nouvelles recherches sur la distribution. *Bulletin de la Société Vaudoise des Sciences Naturelles*, 44, 223–270.
- Jeh, G., & Widom, J. (2002). SimRank: A measure of structural-context similarity. In *Proceedings of the eighth ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 538–543). Association for Computing Machinery.
- Johar, G. V., & Pham, M. T. (1999). Relatedness, prominence, and constructive sponsor identification. *Journal of Marketing Research*, 36(3), 299–312.
- Johar, G. V., Pham, M. T., & Wakefield, K. L. (2006). How event sponsors are really identified: A (baseball) field analysis. *Journal of Advertising Research*, 46(2), 183–198.
- John, D. R., Loken, B., Kim, K., & Monga, A. B. (2006). Brand concept maps: A methodology for identifying brand association networks. *Journal of Marketing Research*, 43(4), 549–563.
- Kahr, M., Leitner, M., Ruthmair, M., & Sinnl, M. (2021). Benders decomposition for competitive influence maximization in (social) networks. *Omega*, 100, Article 102264.
- Kamins, M. A., & Gupta, K. (1994). Congruence between spokesperson and product type: A matchup hypothesis perspective. *Psychology & Marketing*, 11(6), 569–586.
- Kang, J. H., & Lerman, K. (2012). Using lists to measure homophily on Twitter. In *Workshops at the Twenty-Sixth AAAI Conference on Artificial Intelligence July 22–23*.
- Kaur, K. (2016). Development of a framework for analyzing terrorism actions via Twitter lists. In *2016 International conference on computational techniques in information and communication technologies (ICCTICT)* (pp. 19–24). IEEE.
- Keller, K. L. (2013). *Strategic brand management: Building, measuring, and managing brand equity* (4th ed.). Hoboken, NJ: Prentice Hall.
- Kietzmann, J., Paschen, J., & Treen, E. (2018). Artificial intelligence in advertising: How marketers can leverage artificial intelligence along the consumer journey. *Journal of Advertising Research*, 58(3), 263–267.
- Kim, D., Jo, Y., Moon, I. C., & Oh, A. (2010). Analysis of Twitter lists as a potential source for discovering latent characteristics of users. ACM CHI workshop on microblogging (Vol. 6). New York: Association for Computing Machinery.
- Knittel, C. R., & Stango, V. (2014). Celebrity endorsements, firm value, and reputation risk: Evidence from the Tiger Woods scandal. *Management Science*, 60(1), 21–37.
- Knoll, J., & Matthes, J. (2017). The effectiveness of celebrity endorsements: A meta-analysis. *Journal of the Academy of Marketing Science*, 45(1), 55–75.
- Kommiya Mothilal, R., Mishra, D., Nishal, S., Lalani, F. M., & Pal, J. (2022). Voting with the stars: Analyzing partisan engagement between celebrities and politicians in India. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW1), 1–29.
- Kraut, R., Burke, M., Riedl, J., & Resnick, P. (2010). *Dealing with newcomers*. evidence-based social design: Mining the social sciences to build online communities. Cambridge, MA: MIT Press.
- Lee, T. Y., & Bradlow, E. T. (2011). Automated marketing research using online customer reviews. *Journal of Marketing Research*, 48(5), 881–894.
- Levinthal, B. R., & Franconeri, S. L. (2011). Common-fate grouping as feature selection. *Psychological Science*, 22(9), 1132–1137.
- Li, H., Gupta, A., Zhang, J., & Flor, N. (2020). Who will use augmented reality? An integrated approach based on text analytics and field survey. *European Journal of Operational Research*, 281(3), 502–516.
- Ludwig, S., De Ruyter, K., Friedman, M., Brüggem, E. C., Wetzels, M., & Pfann, G. (2013). More than words: The influence of affective content and linguistic style matches in online reviews on conversion rates. *Journal of Marketing*, 77(1), 87–103.
- Maheswaran, D., & Sternthal, B. (1990). The effects of knowledge, motivation, and type of message on ad processing and product judgments. *Journal of Consumer Research*, 17(1), 66–73.
- Mai, L., & Le, B. (2021). Joint sentence and aspect-level sentiment analysis of product comments. *Annals of Operations Research*, 300(2), 493–513.
- Makrides, A., Vrontis, D., & Christofi, M. (2020). The gold rush of digital marketing: Assessing prospects of building brand awareness overseas. *Business Perspectives and Research*, 8(1), 4–20.
- Malhotra, P., & Bhattacharyya, S. (2022). EXPRESS: Leveraging co-followerhip patterns on social media to identify brand alliance opportunities. *Journal of Marketing*. <https://doi.org/10.1177/00222429221083668>.
- Matherly, T. (2019). A panel for lemons? Positivity bias, reputation systems and data quality on MTurk. *European Journal of Marketing*, 53(2), 195–223.
- Meisel, S., & Mattfeld, D. (2010). Synergies of operations research and data mining. *European Journal of Operational Research*, 206(1), 1–10.
- McDaniel, S. R. (1999). An investigation of match-up effects in sport sponsorship advertising: The implications of consumer advertising schemas. *Psychology & Marketing*, 16(2), 163–184.
- Miciak, A. R., & Shanklin, W. L. (1994). Choosing celebrity endorsers. *Marketing Management*, 3(3), 50–59.
- Misra, S., & Beatty, S. E. (1990). Celebrity spokesperson and brand congruence: An assessment of recall and affect. *Journal of Business Research*, 21(2), 159–173.
- Molina-González, M. D., Martínez-Cámara, E., Martín-Valdivia, M. T., & Perea-Ortega, J. M. (2013). Semantic orientation for polarity classification in Spanish reviews. *Expert Systems with Applications*, 40(18), 7250–7257.
- Moore, V. (2007). Where did it all go wrong for the beer they call “wife beater”? *Daily Mail*. <https://www.dailymail.co.uk/news/article-494149/Where-did-wrong-beer-wife-beater.html>. (retrieved September 18, 2019).
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92(3), 289–319.
- Nam, H., Joshi, Y. V., & Kannan, P. K. (2017). Harvesting brand information from social tags. *Journal of Marketing*, 81(4), 88–108.
- Nam, H., & Kannan, P. K. (2014). The informational value of social tagging networks. *Journal of Marketing*, 78(4), 21–40.
- Nasirifard, P., & Hayes, C. (2011). Tadvice: A Twitter assistant based on Twitter lists. In *International conference on social informatics* (pp. 153–160). Springer.
- 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.
- Ng, S., & Houston, M. J. (2009). Field dependency and brand cognitive structures. *Journal of Marketing Research*, 46(2), 279–292.
- Nielsen, F.Å. (2011). A new ANEW: Evaluation of a word list for sentiment analysis in microblogs. arXiv preprint arXiv:1103.2903.
- Oikonomidou, M., Pratikakis, P., Saridakis, C., & Angelidou, S. (2019). Mining celebrity endorsement perceptions using varieties of Twitter account automated data. In *Proceedings of the 6th ACM conference, celebration of women in computing: WomEncourage*. Association for Computing Machinery.
- Olafsson, S., Li, X., & Wu, S. (2008). Operations research and data mining. *European Journal of Operational Research*, 187(3), 1429–1448.
- Olson, E. L., & Thjomøe, H. M. (2009). Sponsorship effect metric: Assessing the financial value of sponsoring by comparisons to television advertising. *Journal of the Academy of Marketing Science*, 37(4), 504–515.
- Page, L., Brin, S., Motwani, R., & Winograd, T. (1999). *The pagerank citation ranking: Bringing order to the web*. Stanford InfoLab.
- Park, C. W., Milberg, S., & Lawson, R. (1991). Evaluation of brand extensions: The role of product feature similarity and brand concept consistency. *Journal of Consumer Research*, 18(2), 185–193.
- Parmar, Y., Ghuman, M. K., & Mann, B. J. S. (2020). The match-up between celebrity associations and product type. *Journal of Creative Communications*, 15(1), 65–89.
- Paul, I., Khattar, A., Chopra, S., Kumaraguru, P., & Gupta, M. (2019). What sets verified users apart? Insights, analysis and prediction of verified users on Twitter. In *Proceedings of the 10th ACM conference on web science* (pp. 215–224). Association for Computing Machinery.
- Pauwels, K., & Van Ewijk, B. (2013). Do online behavior tracking or attitude survey metrics drive brand sales? An integrative model of attitudes and actions on the consumer boulevard. *Marketing Science Institute Working Paper Series*, 13(118), 1–49.
- Peng, J., Agarwal, A., Hosanagar, K., & Iyengar, R. (2018). Network overlap and content sharing on social media platforms. *Journal of Marketing Research*, 55(4), 571–585.
- Pishko, J. (2019). *When brands tease their rivals on Twitter: Consequences for consumer engagement*. University of Houston Doctoral Dissertation.
- Puranam, D., Narayan, V., & Kadiyali, V. (2017). The effect of calorie posting regulation on consumer opinion: A flexible latent Dirichlet allocation model with informative priors. *Marketing Science*, 36(5), 726–746.
- Rakesh, V., Singh, D., Vinzamuri, B., & Reddy, C. K. (2014). Personalized recommendation of Twitter lists using content and network information. In *Proceedings of the eighth international AAAI conference on weblogs and social media*. AAAI Press.
- Rifon, N. J., Choi, S. M., Trimble, C. S., & Li, H. (2004). Congruence effects in sponsorship: The mediating role of sponsor credibility and consumer attributions of sponsor motive. *Journal of Advertising*, 33(1), 30–42.
- Rigby, R. (2011). Brand damage by association: When the wrong type of customer buys your product. *Financial Times*. <https://www.ft.com/content/66d1d3ec-2f19-11e0-88ec-00144feabdc0>. (retrieved July 28, 2019).
- Ringel, D. M., & Skiera, B. (2016). Visualizing asymmetric competition among more than 1,000 products using big search data. *Marketing Science*, 35(3), 511–534.
- Rodgers, S. (2003). The effects of sponsor relevance on consumer reactions to inter-brand sponsorships. *Journal of Advertising*, 32(4), 67–76.
- Rosch, E. (1978). Principles of categorization. In E. Rosch, & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.
- Rosen, S. (1983). The economics of superstars: Reply. *The American Economic Review*, 73(3), 460–462.
- Rozin, P., & Kalat, J. W. (1971). Specific hungers and poison avoidance as adaptive specializations of learning. *Psychological Review*, 78(6), 459–486.
- Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, 5(4), 296–320.
- Ruths, D., & Pfeffer, J. (2014). Social media for large studies of behavior. *Science*, 346(6213), 1063–1064.
- Scholz, M., Pfeiffer, J., & Rothlauf, F. (2017). Using PageRank for non-personalized default rankings in dynamic markets. *European Journal of Operational Research*, 260(1), 388–401.
- Schweidel, D. A., & Moe, W. W. (2014). Listening in on social media: A joint model of sentiment and venue format choice. *Journal of Marketing Research*, 51(4), 387–402.
- Shama, A., & Dhage, S. N. (2018). Application of associative network theory to mine relevant aspect terms from customer reviews. In *2018 International conference on current trends towards converging technologies (ICCTCT)* (pp. 1–7). IEEE.

- Shapiro, S., Spence, M. T., & Gregan-Paxton, J. (2009). Factors affecting the acquisition and transfer of novel attribute relationships to new product categories. *Psychology & Marketing*, 26(2), 122–144.
- Sisario, B. (2012). Beyoncé Deal, Pepsi focuses on collaboration. *The New York Times*. <https://www.nytimes.com/2012/12/10/business/media/in-beyonce-deal-pepsi-focuses-on-collaboration.html>. (retrieved November 3, 2019).
- Smith, S. M., Roster, C. A., Golden, L. L., & Albaum, G. S. (2016). A multi-group analysis of online survey respondent data quality: Comparing a regular USA consumer panel to MTurk samples. *Journal of Business Research*, 69(8), 3139–3148.
- Spry, A., Pappu, R., & Cornwell, T. B. (2011). Celebrity endorsement, brand credibility and brand equity. *European Journal of Marketing*, 45(6), 882–908.
- Statista (2020). Distribution of Twitter users worldwide as of July 2020, by gender. <https://www.statista.com/statistics/828092/distribution-of-users-on-twitter-worldwide-gender/>(retrieved September 2, 2020).
- Steenkamp, J. B., & Van Trijp, H. (1997). Attribute elicitation in marketing research: A comparison of three procedures. *Marketing Letters*, 8(2), 153–165.
- Sul, H. K., Dennis, A. R., & Yuan, L. (2017). Trading on Twitter: Using social media sentiment to predict stock returns. *Decision Sciences*, 48(3), 454–488.
- Sumathy, K. L., & Chidambaram, M. (2013). Text mining: Concepts, applications, tools and issues—an overview. *International Journal of Computer Applications*, 80(4), 29–32.
- Symitsi, E., Stamolampros, P., Daskalakis, G., & Korfiatis, N. (2021). The informational value of employee online reviews. *European Journal of Operational Research*, 288(2), 605–619.
- Teichert, T. A., & Schöntag, K. (2010). Exploring consumer knowledge structures using associative network analysis. *Psychology & Marketing*, 27(4), 369–398.
- Tian, S., Tao, W., Hong, C., & Tsai, W. H. S. (2021). Meaning transfer in celebrity endorsement and co-branding: Meaning valence, association type, and brand awareness. *International Journal of Advertising*, 1–21.
- Tirunillai, S., & Tellis, G. J. (2014). Mining marketing meaning from online chatter: Strategic brand analysis of big data using latent Dirichlet allocation. *Journal of Marketing Research*, 51(4), 463–479.
- Tomikawa, S. A., & Dodd, D. H. (1980). Early word meanings: Perceptually or functionally based? *Child Development*, 51(4), 1103–1109.
- Torres, A., & Bijmolt, T. H. (2009). Assessing brand image through communalities and asymmetries in brand-to-attribute and attribute-to-brand associations. *European Journal of Operational Research*, 195(2), 628–640.
- Urban, G. L., Johnson, P. L., & Hauser, J. R. (1984). Testing competitive market structures. *Marketing Science*, 3(2), 83–112.
- U.S. Census Bureau (2019). Quick facts. <https://www.census.gov/newsroom/press-releases/2019/estimates-characteristics.html> (retrieved September 5, 2020).
- Velichety, S., & Ram, S. (2013). Examining lists on twitter to uncover relationships between following, membership and subscription. In *Proceedings of the 22nd international conference on world wide web* (pp. 673–676). Association for Computing Machinery.
- Wakabayashi, D. (2018). Inside the mating rituals of brands and online stars. *The New York Times*. <https://www.nytimes.com/2018/07/15/technology/online-stars-brands.html>. (retrieved January 15, 2019).
- Wright, O. (2012). Lobbying company tried to wipe out “wife beater” beer references. *Independent*. <https://www.independent.co.uk/tech/lobbying-company-tried-to-wipe-out-wife-beater-beer-references-6284622.html>. (retrieved July 18, 2022).
- Wu, S., Hofman, J. M., Mason, W. A., & Watts, D. J. (2011). Who says what to whom on twitter. In *Proceedings of the 20th international conference on world wide web* (pp. 705–714). Association for Computing Machinery.
- Yamaguchi, Y., Amagasa, T., & Kitagawa, H. (2011). Tag-based user topic discovery using Twitter lists. In *2011 International conference on advances in social networks analysis and mining* (pp. 13–20). IEEE.
- Yu, S., & Hu, Y. (2020). When luxury brands meet China: The effect of localized celebrity endorsements in social media marketing. *Journal of Retailing and Consumer Services*, 54, Article 102010.
- Zhan, Y., & Tan, K. H. (2020). An analytic infrastructure for harvesting big data to enhance supply chain performance. *European Journal of Operational Research*, 281(3), 559–574.
- Zhao, J., & Ram, S. (2011). Examining the evolution of networks based on lists in Twitter. In *2011 IEEE 5th international conference on internet multimedia systems architecture and application* (pp. 1–8).
- Zhao, Y., Yang, S., Narayan, V., & Zhao, Y. (2013). Modeling consumer learning from online product reviews. *Marketing Science*, 32(1), 153–169.
- Zheng, Y., & Ni, W. (2020). Analysis on transformation and development of traditional enterprise based on internet celebrity economy. *Journal of Education and Praxis Research*, 1(6), 1–9.