Helpfulness Assessment of Online Reviews: The Role of Semantic Hierarchy of Product Features

YIN KANG, University of Maryland Baltimore County
LINA ZHOU, The University of North Carolina at Charlotte

Effective use of online consumer reviews is hampered by uncertainty about their helpfulness. Despite a growing body of knowledge on indicators of review helpfulness, previous studies have overlooked rich semantic information embedded in review content. Following design science principles, this study introduces a semantic hierarchy of product features by probing the review text. Using the hierarchical framework as a guide, we develop a research model of review helpfulness assessment. In the model, we propose and conceptualize three new factors—breadth, depth, and redundancy, by building on and/or extending product uncertainty, information quality, signaling, and encoding variability theories. The model-testing results lend strong support to the proposed effects of those factors on review helpfulness. They also reveal interesting differences in the effects of redundancy and readability between different types of products. This study embodies knowledge moments of multiple genres of inquiry in design science research, which have multifold research and practical implications.

CCS Concepts: • **Information systems** → Data analytics;

Additional Key Words and Phrases: Review helpfulness, product feature, semantic hierarchy, breadth, depth, redundancy, design science

ACM Reference format:
Yin Kang and Lina Zhou. 2019. Helpfulness Assessment of Online Reviews: The Role of Semantic Hierarchy of Product Features. *ACM Trans. Manage. Inf. Syst.* 10, 3, Article 12 (November 2019), 18 pages.
https://doi.org/10.1145/3365538

1 INTRODUCTION

Consumers often read online reviews to seek information on a product, and reviews can help reduce purchasing uncertainty [Dimoka et al. 2012; Markopoulos and Clemons 2013] and diminish search cost and switching cost [Li et al. 2011]. Consumer reviews can further impact product sales [Chevalier and Mayzlin 2006; Zhu and Zhang 2010] through influencing consumer purchase intention [Clemons et al. 2006] and repeated purchasing probability [Li et al. 2011]. Compared with expert reviews, consumer reviews are able to elicit higher perceived usefulness, trusting beliefs, and perceived affective quality for experience goods [Benlian et al. 2012]. However,

This work is supported in part by the National Science Foundation (Awards No. CNS 1704800 and No. SES 1527684). Any opinions, findings or recommendations expressed here are those of the authors and are not necessarily those of the sponsor of this research.
Authors' addresses: Y. Kang, University of Maryland Baltimore County, 1000 Hilltop Cir, Baltimore, MD 21250, University of Maryland Baltimore County, Baltimore, MD, USA; email: ky6@umbc.edu; L. Zhou, The University of North Carolina at Charlotte, Charlotte, NC, USA; email: lzhou8@uncc.edu.
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© 2019 Association for Computing Machinery.
2158-656X/2019/11-ART12 $15.00
https://doi.org/10.1145/3365538

an explosive accumulation of online reviews along with variations in the review quality compromise the value of online reviews to prospective consumers. Assessing review helpfulness to enable the sorting of online reviews based on helpfulness can be a promising way of addressing the above problem.

Helpful reviews can potentially provide product diagnosticity across different stages of consumer purchase decision-making process [Pavlou et al. 2007]. Online retailers that provide more helpful reviews offer greater potential value to their customers [Mudambi and Schuff 2010]. However, many websites do not yet provide the function of helpfulness voting. Even when the function is available, most reviews only receive less than a handful of and even no helpfulness votes. More importantly, there is generally lack of explanation of the rationale behind the helpfulness voting beyond a simple total count vote. Consequently, it limits the role of online reviews in support of making purchase decisions.

Despite a growing body of literature on review helpfulness (e.g., Filieri et al. [2018], Ghose and Ipeirotis [2011], Hong et al. [2017], Mudambi and Schuff [2010], and Yin et al. [2016]), they have mainly relied on star rating and word-count-based statistics of review text. It has been widely recognized that numerical rating is insufficient, because it does not provide justifications or reasoning for the rating [McAuley and Leskovec 2013]. For instance, a product, despite low quality, may receive an overall four-star rating due to its low price. Wu et al. [2015] argue that textual comments accompanying numerical ratings are more valuable than ratings themselves. Review text has been explored for review helpfulness (e.g., Hong et al. [2017] and Yin et al. [2014]). Nonetheless, most of them characterize review text by variables defined on the basis of word count, such as review length. While longer reviews tend to provide more details about focal products, they tend to be more difficult to comprehend [Kuan et al. 2015]. Review length also provides contradictory evidence (see also Kang and Zhou [2016]) and thus is insufficient for explaining review helpfulness.

Review text contains rich information that can be leveraged to uncover latent product dimensions, among others [McAuley and Leskovec 2013]. Product features comprises one key aspect of products. However, only a handful of studies have explored the influence of product features on review helpfulness to date (e.g., Archak et al. [2011], Kim et al. [2006], and Li et al. [2013]), and these studies are still limited in that: (1) they treat different product features as orthogonal without considering their semantic relationships; (2) they assume that all product features are equally important for helpfulness assessment; (3) they do not provide quantifiable measures of product features beyond simple count; and (4) their methods have limited scalability due to their reliance on human judgment.

This research involves both build-and-evaluate and justify-and-theorize activities [Hevner et al. 2004; Walls et al. 1992]. It follows the "design and evaluate" cycle in the design and development of a new research model for review helpfulness assessment. Specifically, we first introduce product feature hierarchy as a framework, develop approaches to hierarchy construction, and evaluate the results of hierarchy construction in Section 3; then we propose a model of review helpfulness assessment by applying the semantic hierarchy and drawing on multiple theories, such as product uncertainty, information quality, signaling, information processing, and encoding variability theories in Section 4, and we evaluate the model with regression analysis and robustness tests on real online review data in Section 5; finally, we demonstrate practical impacts of the research model by developing predictive models that incorporate the three novel factors—breadth, depth, and redundancy.

2 BACKGROUND AND RELATED WORK

We first define review helpfulness by drawing on product uncertainty theory [Dimoka et al. 2012], which accounts for both expectation and subjective evaluation of products.

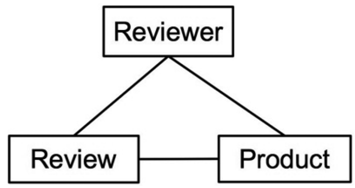

Fig. 1. A schema of online review evaluation.

Definition 1. Review helpfulness is the extent to which a review can help consumers reduce product uncertainty in determining the true quality of a product and achieving a fit between consumers' expectations and the reviewer's evaluation.

Based on the above definition, the helpfulness of reviews is centered on their ability to reduce product uncertainty by providing information that is diagnostic to consumer purchase decisions. Accordingly, the indicators of review helpfulness can be approached along three dimensions: product, review, and reviewer. To this end, we propose a three-dimensional schema of online review evaluation (see Figure 1). Each of the dimensions in turn consists of a set of indicators. For example, product features constitute a key indicator of the product dimension; length, readability, and rating are representatives of the review dimension; and reviewer reputation and social interactions are indicators of the reviewer dimension. The schema serves as guidance in organizing the related work, as summarized in Table 1.

Compared with other dimensions, the product dimension has received the least attention to date. Studies suggest that presenting more information about products improve diagnositicity and reduce product uncertainty, which in turn make reviews helpful [Mudambi and Schuff 2010]. Product features represent detailed information about products. Among the limited studies on review helpfulness that have used product features (e.g., Archak et al. [2011]), they have treated different product features as orthogonal without considering their semantic relationships. In reality, some pairs of features are semantically more closely related to each other than other pairs. For example, *battery life* is more closely related to *battery*, and *battery weight* is more closely related to *battery life*, than *lens*, of a camera. In addition, different text forms may be used to express the same product features (e.g., *picture* and *photo*). As a result, ignoring the semantic relationships may lead to failure in delineation among product features in terms of their specificity and coverage. Therefore, accounting for the semantic relationships between product features in review texts can be instrumental to reducing product uncertainty in consumer purchase decision-making.

3 SEMANTIC HIERARCHY OF PRODUCT FEATURES AND ITS CONSTRUCTION

3.1 A Semantic Hierarchy of Product Features

We extend the scope of product features [Liu 2011] by explicitly incorporating service features pertaining to product transactions and meta-features of product features.

Definition 2. Product features refer to attributes, parts, functions, and related concepts and services associated with transactions of a product, and of a part, subpart, or specific case of functions.

Product feature is recursive by definition, which motivates us to propose a semantic hierarchy of product features. In the hierarchy, every product feature, except for the root feature, is connected to another one via a subsumption relation (e.g., *is-a* and *has-a*). As shown in Figure 2(a), features of a product (level 1) can come in different types such as attributes, parts, and functions (level 2), and accordingly the latter are subsumed under the former layer; the level 2 product features can in turn subsume other product features (e.g., subparts, sub-attribute, and sub-functions) (level 3), and so

Table 1. A Summary of Indicators of Review Helpfulness

Dimension	Indicator	Description and Sample Studies
Product	Product feature	Number of product features [Kim et al. 2006]
		Aggregation of product features [Archak et al. 2011]
Review	Review length	Word counts [Chen and Tseng 2011; Kang and Zhou 2016; Mudambi and Schuff 2010] [Hong et al. 2017; Kuan et al. 2015; Singh et al. 2017]
	Readability	Easy to understand [Hong et al. 2017; Hu et al. 2012; Karimi and Wang 2017; Kuan et al. 2015; Schindler Robert and Bickart 2012; Singh et al. 2017]
	Writing style	The use of inexpressive slang, qualifications, repetition, humor, formal language, expressions of emotion, and first-person pronouns [Li et al. 2013; Malik and Hussain 2018; Schindler Robert and Bickart 2012]
	Rating	Star rating [Chua and Banerjee 2016; Danescu-Niculescu-Mizil et al. 2009; Mudambi and Schuff 2010; Singh et al. 2017; Yin et al. 2014]
	Sentiment	Emotions, sentiments [Kuan et al. 2015; Ren and Hong 2018; Ullah et al. 2016; Yin et al. 2014]
	Review order	Order of review posting [Zhou and Guo 2017]
Reviewer	Reputation	Reviewer disclosure and reviewer history [Filieri et al. 2018; Ghose and Ipeirotis 2011; Hong et al. 2017; Jensen et al. 2013; Kuan et al. 2015; Malik and Hussain 2018; Weathers et al. 2015]
	Social interaction	Reference to other reviews; online social networking, online product referral [O'Mahony and Smyth 2010; Zheng et al. 2011; Zhou and Guo 2017]
	Reviewer Profile	Visual cues, reviewer profile image [Karimi and Wang 2017]

on. In other words, the hierarchy can grow deeper by including even finer-grained product features via the subsumption relation. Figures 2(b) and 2(c) are illustrations of the semantic hierarchy with two products: camera and movie, respectively. For instance, *lens* is a part of a *camera*, and *cap* is a part, and *size* is an attribute of lens. Similarly, *action* is an attribute of movie, and *fight* and *car-chasing* are parts of *action*. The semantic hierarchy can be defined at the product level (product hierarchy) or at the review level (review hierarchy).

Definition 3. Product hierarchy is composed of F, the entire set of features of a focal product, and subsumption relations existing between f_i and f_j, where $f_i \sqsubset_T f_j$ and $f_i, f_j \in F$.

Definition 4. Review hierarchy of review r is a sub-hierarchy of the corresponding product hierarchy derived from the mentioned features $\{f | f \in F_r, \ F_r \subseteq F\}$, where F_r is a set of product features that are mentioned in r. It is composed of F_r and other product features that are directly or indirectly connected to f via subsumption relations in the product hierarchy.

The framework can serve as a guide in the development of our research model for two main reasons: (1) it provides a fine-grained representation of a product in terms of its features, which allows consumers to gain an in-depth understanding of online reviews; and (2) it offers a systematic

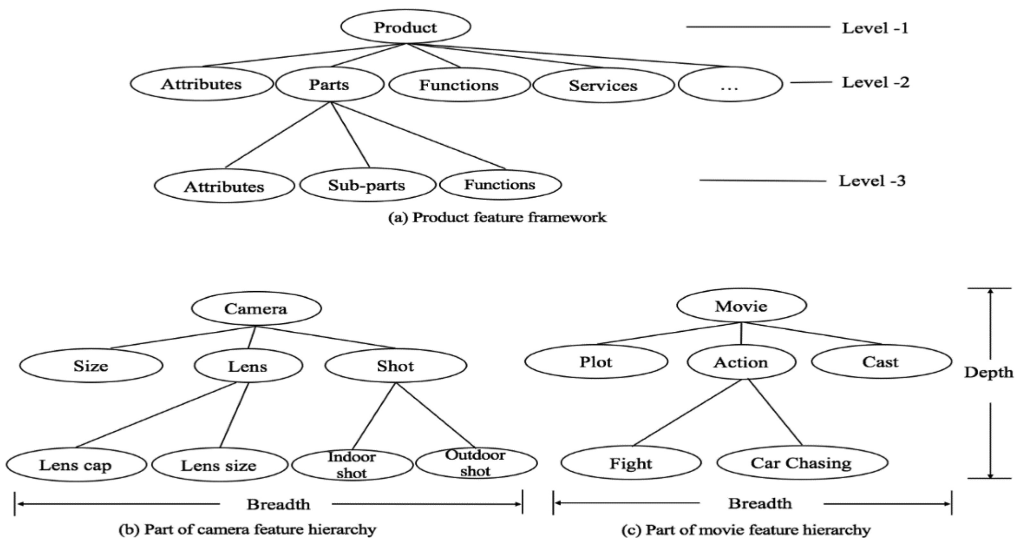

Fig. 2. A hierarchical framework of product features and two illustrations.

way of identifying potentially influencing factors in various dimensions (e.g., breadth and depth). A given product can receive any number of reviews. Accordingly, we can derive a review hierarchy for each of the reviews from the product hierarchy.

Given the backdrop of review hierarchy, we propose three main variables of online reviews: breadth, depth, and redundancy. First, breadth and depth are key properties of a hierarchical tree structure, which can also be used to characterize a review hierarchy in terms of its coverage and specificity of product features, respectively. Second, in view that two review hierarchies of the same product can be similar to each other by sharing some common product features, we introduce redundancy as a third variable of online reviews.

Definition 5. Breadth of review *r* is defined as the number of product feature nodes at a selected level of the review hierarchy of *r*.

Definition 6. Depth of review *r* is defined as the average level of mentioned product features in review hierarchy of *r*. The level of a mentioned feature node is defined as the number of subsumption relations that connect the feature to the root node in the hierarchy.

Definition 7. Redundancy of review *r* is defined as the average redundancy of all mentioned product features in review hierarchy of *r*. The redundancy of a mentioned feature node is defined as the number of other reviews of the same product that mention the same product feature in their review hierarchies.

3.2 Hierarchy Construction

Product feature extraction is instrumental to automatic construction of a semantic hierarchy of product features. Many attempts have been made to extract product features from online consumer reviews. For example, double propagation is designed to extract both product features and opinions based on their dependency relationships [Qiu et al. 2011]; and RubE [Kang and Zhou 2017] extracts both subjective and objective features (associated and not associated with an opinion) by leveraging both direct and indirect dependency relations. We adopt RubE [Kang and Zhou 2017] in this study, because it represents the state-of-the-art performance in product feature extraction, and it is capable to extract both frequent and infrequent features.

Given a set of extracted product features, we perform hierarchy construction by determining subsumption relations (e.g., *is-a* and *has-a*) and relatedness (e.g., semantic similarity) between those features. The method design for hierarchy construction is customized to the characteristics of individual products. For those products that do not yet have a feature taxonomy/ontology or tend to be fast-evolving (e.g., electronics), we adopt a bottom-up approach [Kang et al. 2014]. Starting with a set of extracted features, the approach proceeds in the following steps: (1) grouping similar features based on their relatedness, and each of these groups forms a leaf node in the semantic hierarchy; (2) identifying a parent node (i.e., head feature) from the features within each new node based on subsumption relations, which is promoted to one level higher; (3) grouping parent nodes based on the relatedness of their features; and (4) repeating steps 2~3 until all current top-level nodes are merged into a single root node (i.e., product).

For other products that have taxonomy or ontology of product features available (e.g., movie), we employ a middle-out approach. The approach starts with identifying a set of key features, and then uses them to derive both parent-node (i.e., more general) features and child-node (i.e., more specific) features simultaneously. The latter step can be repeated to further grow the hierarchy. For instance, drawing concepts from a movie review ontology [Zhou and Chaovalit 2008], we first create key movie features as the base set, and then we explore possible subsumption relations between base set features and extracted candidate features from online reviews to construct a semantic hierarchy. The second task is performed independently by two human coders who are familiar with movie reviews. The resulting product feature hierarchies achieved an acceptable level of consistency on the identified movie feature pairs based on an inter-rater agreement statistic (kappa coefficient $\kappa = 0.761$). We reconcile the inconsistent results via face-to-face meetings with the coders. Once a product hierarchy is constructed, we can derive review hierarchies for individual reviews for the same product.

The semantic hierarchy of product features enables a fine-grained feature-based representation of online reviews. The key properties of the semantic hierarchy lend themselves to identifying important concepts for deep understanding of review helpfulness.

4 THEORETICAL FOUNDATION AND HYPOTHESES DEVELOPMENT

4.1 Theoretical Foundation

Consumers often have to make decisions with incomplete information about products, sellers, and the available alternatives, consequently they are forced to face a variety of uncertainties [Dimoka et al. 2012]. Product uncertainty refers to the consumer's difficulty in evaluating product attributes and predicting how a product will perform in the future, which is a major impediment to online markets [Dimoka et al. 2012]. As a major dimension of product uncertainty, product quality uncertainty deals with vertically differentiated product features that offer common utility to consumers [Garvin 1984], and can be mitigated by exposure to quality information such as information about product features [Markopoulos and Clemons 2013].

Based on the information signaling literature [Pavlou et al. 2007], product information signals embedded in a consumer review is one major approach to mitigate product uncertainty explicitly. Information signals help consumers infer the value of products with unobservable quality and uncertain value [Crawford and Sobel 1982], and facilitate their decision making [Urbany et al. 1989]. In addition, consumer reviews are often considered as more credible and trustworthy than expert reviews [Benlian et al. 2012], filling the information gap caused by product uncertainty. Reading online consumer reviews can be viewed as a process of increasing product familiarity and product quality diagnositicity, and gaining third party assurance; and accordingly, helpfulness evaluation is viewed as a process of communicating with prospective consumers about whether a review meets

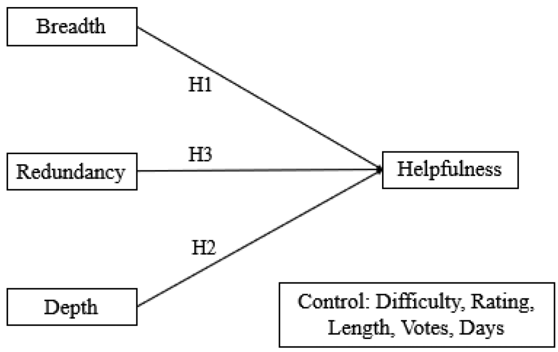

Fig. 3. The Research Model.

the above expectations. In other words, online reviews help reduce description and performance of product uncertainty by providing information, which are expressed as comments, opinions, or experiences on product features. Therefore, it is natural and essential to integrate product uncertainty theory and information quality theory in evaluating the information expressed in online reviews against product uncertainty reduction.

Effective information signals must be visible, clear, credible, and differentially costly [Rao and Monroe 1989]. The diagnosticity of textual product descriptions is proposed to be an effective signal that can help buyers reduce both their description uncertainty by giving detailed information on the product's characteristics and performance uncertainty by helping buyers infer how the product will perform in the future based on information on its current condition, usage, and experience [Pavlou et al. 2007].

Information quality (IQ) has been used to explain review helpfulness [Chen and Tseng 2011; Chua and Banerjee 2016]. Nonetheless, the studies neither make an attempt at theory development, nor examines information quality in light of product uncertainty reduction while evaluating online reviews. IQ is a multidimensional construct, with each dimension capturing one aspect of information that is important to its consumers [Wang and Strong 1996]. Unlike the previous studies that either select IQ dimensions in an ad hoc manner [Chua and Banerjee 2016] or mix together all constructs of an IQ framework [Chen and Tseng 2011], we focus on contextual IQ [Wang and Strong 1996] due to that it captures quality dimensions from the perspective of information consumers and attempts to strike a balance between theoretical consistency and practicability [Ghasemaghaei and Hassanein 2015]. Contextual IQ is composed of six dimensions including completeness, amount of information, relevancy, value-added, and timeliness [Wang and Strong 1996]. It emphasizes that quality of information must be assessed in the context of the task being performed.

In this research, we develop a research model of review helpfulness assessment by drawing on the above and other related theories (see Figure 3). Specifically, we use information quality theory mainly to conceptualize new factors or IVs of the proposed research model, product uncertainty and encoding variability to hypothesize the effects of IVs on DV (review helpfulness), signaling theory to argue for the role of information contained in online reviews in reducing product uncertainty, and information processing theory to connect the information contained in review texts and prospective consumers.

4.2 Hypotheses Development

Uncertainty reduction theory [Berger and Calabrese 1975] asserts that, people need information about the other party to reduce their uncertainty during the initial interaction. Additionally, communication is the main approach of uncertainty reduction. Product uncertainty is rooted in the problem of information asymmetry [Dimoka et al. 2012]. Reading a review can be considered as a process of communication for consumers to collect information of a product for reducing product quality uncertainty. Thus, review helpfulness assessment serves as an alternative way to address the information asymmetry by determining the degree to which a review helps reduce the uncertainty in consumer purchase decisions. The semantic hierarchy of product features in turn supports information quality assessment of online reviews in terms of uncertainty reduction by probing into the semantic relationships of product features embedded in review texts.

Product diagnosticity is one of the major mitigating factors to overcome the barrier of perceived information asymmetry in e-commerce [Pavlou et al. 2007]. It has been argued that a helpful review must provide complete information [Chen and Tseng 2011], and the quantity or volume of information embedded in a review can help reduce product uncertainty [Kang and Zhou 2016]. Complete information contributes to product diagnosticity by better facilitating buyers in the evaluation of product quality and seller's true characteristics [Kumar and Benbasat 2006].

We use breadth to represent the degree of information completeness or amount of information expressed in an online review. Review text has the potential to cover a wide range of information about a product, such as pacing, cast, and plot of a movie [Zhou and Chaovalit 2008]. Completeness and amount of information are key dimensions of contextual IQ [Wang and Strong 1996]. The breadth of a review hierarchy can reflect both of these dimensions. The hierarchical framework of product features helps conceptualize and operationalize information completeness. For a selected level (e.g., level-2) in a hierarchy of product features (see Figure 2), the more feature nodes that a review hierarchy covers, the greater the breadth of the hierarchy, and the more complete and greater amount of information that the review conveys about a target product. Accordingly, a review with greater breadth is more likely to help other consumers diagnose product quality and make purchase decisions.

One aspect of the information processing theory involves understanding how many discrete units of information can be retained in short-term memory before information loss occurs. Compared with the traditional information unit defined at the word level [Kim et al. 2006], product features can better serve the processing of online reviews, because they are semantically meaningful and self-contained [Kang and Zhou 2016; Miller 1956]. Therefore, we propose the first hypothesis as the following:

H1: Breadth has a positive effect on review helpfulness.

Depth reflects the relevancy and accuracy of information in an online review by measuring the degree of specificity or concreteness of product features mentioned in a review (see Definition 6). In the semantic hierarchy of product features (see Figure 2), a product feature at a lower level (e.g., lens cap or size) is more concrete, specific, and/or precise than its higher-level counterparts (e.g., lens), whereas a higher-level product feature represents a more abstract, general, and/or broad concept than its lower-level counterparts. The information needs of consumers are very specific [Murray 1991], and thus more specific information is expected to be more relevant. The relevancy dimension of contextual IQ measures the extent to which information is applicable and helpful for the task at hand [Wang and Strong 1996]. According to product uncertainty theory [Dimoka et al. 2012], sellers may not be willing or able to provide detailed description about their products. Thus, a review that contains relevant product information can potentially fill the void by helping reduce purchase uncertainty.

In cognitive science [Posner 1989], abstract features are difficult to understand and evaluate, because they cannot be physically perceived. Conversely, concrete consumer reviews are perceived as more diagnostic than abstract ones [Li et al. 2013]. Additionally, concrete or specific reviews are more easily stored in memory because of their distinct boundaries for transmitting information signals to consumers, whereas abstract reviews are the opposite due to their subjective characteristics [Li et al. 2013]. By the same token, the concreteness of information has a critical impact on decision behavior [Borgida and Nisbett 1977]. For example, the use of specific keywords in advertisement can directly improve the sales of advertised products [Lu and Zhao 2014], and users of knowledge management systems are often looking for specific knowledge [Liang et al. 2006]. One of the key factors in effective encoding of information is ensuring that the material is meaningful, and accordingly including concrete and accurate information about product features makes information processing efficient and effective. On a related note, chunking (breaking the information up into manageable chunks) is an important way of assisting encoding information, implying the lower level the information (i.e., a lower-level product feature), the better it supports information processing [Miller 1956]. Furthermore, a consumer review with greater depth contributes to the amount of information by providing more details with less unambiguity to support product quality diagnostics. Therefore, we propose the following hypothesis:

H2: Depth has a positive effect on review helpfulness.

Redundancy arises when other sources only repeat the information from the first source in an alternative way [Schnotz and Kürschner 2007]. Redundancy of a review, which represents the degree to which the same product features are mentioned repeatedly across different reviews (see Definition 7). Online review consumption as a knowledge acquisition process. There is a long history of research into learning via repetition. When applied to advertising repetition and consumer memory, distributed presentation schedule are generally found to be more effective than massed schedule [Sawyer et al. 2003]. Accordingly, it would be able to better support consumer purchase decision making if specific product features are repetitively mentioned across different online reviews than concentrated in a few reviews. Encoding variability theory suggests that presenting a series of ads that contain slight variations on a topic enhances memory for the ad content (see also Unnava and Burnkrant [1991]). Thus, consumers can learn better about a product with repeated inputs on the same product features from different reviews.

From the principal-agent perspective [Pavlou et al. 2007], however, redundant features are unlikely to provide novel and timely information, or reveal hidden information that might contribute to reduced product uncertainty. In addition, information processing theory [Miller 1956] suggests, redundant information cannot be viewed as meaningful and thus would not be stored in limited working memory. Information theory [Shannon 2001] also predicts that repeated features are less valuable, because they have a smaller effect on reducing uncertainty in relation to infrequent features. Therefore, too many repetitions of the same themes or product features would likely minimize the information value of online reviews. We hypothesize that, as the level of redundancy increases, helpfulness initially increases, and at some point, it levels off and then starts to decline.

H3: Redundancy has a curvilinear effect on review helpfulness.

5 METHODS

5.1 Data Collection and Preprocessing

We collect data from Amazon.com, because (1) it has one of the most active communities of online reviewers [Chevalier and Mayzlin 2006]; and (2) it provides helpfulness votes, which can serve as the ground truth for review helpfulness. We select online reviews of electronics (e.g., camera, cell phone, and printer) and movies (Blu-ray), which represent two types of products—search goods and experience goods [Nelson 1970], to test the generality of our proposed model. The differentiation of product types is mainly based on whether their quality can be evaluated by measurement of given

Table 2. Descriptive Statistics of Review Datasets

Product	# of Reviews	# of Features
Camera	2,902	355
Cell phone	2,580	260
Printer	1,430	271
Movie	1,671	155

features [Huang et al. 2013]. Within each product category, we include a wide range of products with respect to brand (e.g., Canon, Nikon), model (point and shoot, DSLR), and price range (e.g., \$50~\$1200).

The review data contains information such as review text, review title, date, and rating. The data goes through the following main preprocessing steps: spam review removal, missing votes filtering, basic natural language processing, feature extraction, and hierarchy construction. Spam activities are widespread and consequential, because extremely positive or negative reviews have a significant effect on consumer purchase intention, resulting in significant financial gain or loss [Lu et al. 2013]. The most common form of spam review is duplicate reviews or near duplicate reviews on the same products, which may cause certain features to appear more frequently. Spam review detection is beyond our scope, so we adopt a simple method - shingle method [Jindal and Liu 2007] using a similarity threshold of 0.9 to remove spam reviews. Additionally, we remove those reviews that have received less than five votes [Kim et al. 2006] to constrain variance (e.g., the amount by which the helpfulness votes would change if we estimated it using a different set of reviewers). The remaining reviews are processed using conventional natural language techniques, including tokenization, POS tagging, and syntactic dependency analysis. Finally, we perform feature extraction and hierarchy construction for each product separately (see Section 3). Descriptive statistics of the datasets is reported in Table 2.

5.2 Operationalization of Research Variables

5.2.1 Dependent and Independent Variables. The dependent variable is review helpfulness, which is operationalized as n/m, meaning that n out of m people found the review helpful.

The independent variables include review breadth, depth, and redundancy. Based on Definitions 5~7, the operationalizations of the variables for review r are presented in Equations (1)~(3):

$$breadth_r = \left|map_2\left(f \mid f \in F_r\right)\right|, \tag{1}$$

$$depth_r = \frac{1}{|F_r|} \sum_{f \in F_r} \left(l\left(f\right) \times c\left(f\right)\right), \tag{2}$$

$$redundancy_r = \frac{1}{|R| \times \left|\{f \in F_r\}\right|} \sum_{\{f \in F_r\}} n\left(f\right), \tag{3}$$

where F_r denotes all product features mentioned in r, map2(·) is a function of mapping · to a level-2 node in review hierarchy of r, |·| denotes the cardinality of set ·; l(·) denotes the level of feature node · in the review hierarchy of r, c(·) the mention frequency of · in r; R denotes the entire set of reviews on a product, and n(·) the number of reviews that mention feature ·.

We choose the second level to operationalize breadth for two main reasons: (1) it is applicable regardless of the depth of a product hierarchy (particularly to a hierarchy with shallow structure); and (2) it helps avoid ambiguities in measuring features in a review hierarchy; otherwise, choosing

Table 3. Descriptive Statistics of all Variables

Variable	Mean	Std.
Helpfulness	0.704	0.295
Breadth	2.865	2.283
Depth	2.203	0.807
Redundancy	0.341	0.258
Length	88.268	83.024
Difficulty	10.826	5.586
Rating	0.865	0.226
Days	526.882	285.759
Votes	37.983	100.931

a lower level (e.g., level 3) would lead to ambiguity in mapping a level-2 product feature (e.g., battery) mentioned in r to a lower-level product feature (e.g., battery weight).

5.2.2 *Control Variables.* We consider the following control variables based on the findings of previous studies.

- Length is measured by word count in the text of review r.
- Rating is measured as star rating normalized by its maximum value (i.e., 5).
- Difficulty (reading difficulty) is measured using Automated Readability Index [Senter and Smith 1967], which approximates the minimal level of education needed to comprehend a text.

To account for the Matthew Effect [Merton 1968], we also include two additional control variables.

- Votes is defined as the total number of votes for review r.
- Days is measured as the number of elapsed days since review r has been posted.

We choose Tobit regression to test the research model (see Equation (4)) in view of the measurement of the dependent variable and the distribution of the sample.

$$\text{Helpfulness} = \alpha_1\text{Breadth} + \alpha_2\text{Depth} + \alpha_3\text{Redundancy} + \alpha_4\text{Redundancy\^{}2} + \alpha_5\text{Length} + \alpha_6\text{Length\^{}2} + \alpha_7\text{Rating} + \alpha_8\text{Rating\^{}2} + \alpha_9\text{Vote} + \alpha_{10}\text{Days} + \alpha_{11}\text{Difficulty} + \varepsilon. \quad (4)$$

We choose two baseline models for comparison. Model 1 incorporates the most widely studied factors in review helpfulness assessment, including rating and length, as both linear and quadratic terms, and vote count; and model 2 is an enhanced version of model 1 that incorporates additional variables from the recent literature, such as reading difficulty, days, and product feature count (*ProdFeatures*).

5.3 Results

The descriptive statistics of all variables is reported in Table 3. Additional tests are conducted to check whether the data are not skewed and whether they are free of the multicollinearity problem. The VIF values of breadth, depth, and redundancy are 2.71, 3.18, and 3.24, respectively. The results confirm that the correlations between those variables and other variables, and among those variables can be neglected.

The results of the regression analyses are reported in Table 4. The table shows that review breadth ($p < 0.001$), depth ($p < 0.01$), redundancy ($p < 0.001$), and redundancy^2 ($p < 0.01$) all

Table 4. Model Test Results on all Reviews

	Model 1	Model 2	Proposed Model
Rating	0.368 (0.023)***	0.256 (0.017)***	0.214 (0.049)***
Rating^2	−0.098 (0.091)**	−0.079 (0.023)*	−0.137 (0.046)**
Length	0.0311 (0.012)**	0.012 (0.093)**	0.035 (0.142)**
Length^2	−0.008 (0.012)**	−0.019 (0.003)†	−0.045 (0.208)†
Votes	0.036 (0.014)***	0.0984 (0.004)**	0.0572 (0.015)***
Difficulty		−0.078 (0.091)†	−0.122 (0.103)†
Days		0.0234 (0.0126)***	0.0191 (0.029)***
ProdFeatures		0.0367 (0.011)**	
Breadth			0.019 (0.049)***
Depth			0.045 (0.027)**
Redundancy			0.216 (0.018)***
Redundancy^2			−0.066 (0.05)**
Adj. R-Squared	0.311	0.335	0.348
AIC	241.27	167.79	148.98
BIC	271.11	222.98	198.74

Notes: Standard errors are shown in parentheses. $^{\dagger}p < 0.1$; $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

have significant effects on review helpfulness. Additionally, length, day, and vote are all found to have positive effects, rating a curvilinear effect, and reading difficulty a negative effect, on review helpfulness, which are consistent with the literature. Therefore, hypotheses H1, H2, and H3 are all supported.

Based on the R-squared statistic reported in Table 4, our proposed model has the best goodness of fit. In addition, we compare our model with the two baseline models using Bayesian Information Criterion (BIC) and Akaike Information Criterion (AIC) [Burnham and Anderson 2002]. The results show that our proposed model yields the lowest BIC and AIC values and thus is the most favored. We further perform a likelihood ratio test, and the results show that our proposed model has a higher log likelihood (−122.56) and hence a better fit than model 1 (−171.15) and model 2 (−136.85).

We validate the practical impacts of the proposed factors on the prediction performance of review helpfulness. Following [Kim et al. 2006], we select SVM regression with the radial basis function kernel as the machine learning technique. In addition, we adopt the Pearson correlation coefficients between the predicted helpfulness scores and actual helpfulness values based on 10-fold cross-validation as the measure of prediction performance. The results show that incorporating the factors introduced in this study as input features improve the correlation coefficient from 0.375 to 0.487.

5.4 Robustness Checks

To increase the confidence in the findings, we conduct a series of robustness tests. First, we consider alternative measures for each of the three independent variables introduced in this study, and their test results are reported in Table 5. The direction and significance of the main effects of breadth, depth, and redundancy remain the same except that the curvilinear effect of redundancy is no longer significant for the alternative measures of depth and breadth.

- *Alternative measure of depth.* Depth of review *r* is operationalized as the most frequent level among all mentioned product features in the review hierarchy of *r*. In case of a tie, the deeper level is selected.

Table 5. Results of the Proposed Model Using Alternative Measures

	Alternative depth	Alternative breadth	Alternative redundancy
Breadth	0.0178 (0.0048)***	0.0129 (0.009)***	0.0355 (0.0091)***
Depth	0.0215 (0.0028)*	0.0066 (0.0034)*	0.0119 (0.0073)**
Redundancy	0.0278 (0.0361)**	0.0604 (0.0921)***	0.0422 (0.051)*
Redundancy^2	−0.0032 (0.0099)	−0.0023 (0.0056)	−0.0038 (0.0009)*
Length	0.0164 (0.0115)**	0.0711 (0.0096)***	0.0156 (0.0011)*
Length^2	−0.0099 (0.0001)	0.0019 (0.0017)	−0.0015 (0.0009)
Rating	0.7516 (0.0547)***	0.3511 (0.0411)**	−0.5681 (0.0342)*
Rating^2	−0.0191 (0.011)	−0.1011 (0.099)	−0.109 (0.0887)
Days	0.0209 (0.0087)**	0.0499 (0.0098)***	0.0501 (0.0077)***
Votes	0.0851 (0.0053)**	0.0726 (0.0011)**	0.0569 (0.018)***
Difficulty	−0.0124 (0.0018)†	−0.0011 (0.0019)	−0.0024 (0.0011)†

Notes: Standard errors are included in parentheses. †$p < 0.1$; *$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

Table 6. Model Test Results on Reviews of two Product Categories

	Search goods ($R^2 = 0.353$)	Experience goods ($R^2 = 0.342$)
Breadth	0.015 (0.002)***	0.013 (0.004)**
Depth	0.029 (0.008)***	0.027 (0.011)*
Redundancy	0.108 (0.007)*	0.003 (0.004)
Redundancy^2	−0.017 (0.001)*	−0.029 (0.012)
Length	0.023 (0.004)***	0.018 (0.006)**
Length^2	−0.002 (0.009)*	0.01 (0.011)
Rating	0.103 (0.004)**	0.061 (0.009)***
Rating^2	−0.091 (0.002)*	−0.009 (0.091)†
Days	0.053 (0.011)**	−0.015 (0.014)*
Votes	0.017 (0.004)***	0.067 (0.008)***
Difficulty	−0.012 (0.004)**	0.001 (0.058)

*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

- *Alternative measure of breadth.* Breadth of review *r* is operationalized as the number of leaf-node features of all mentioned product features in the review hierarchy of *r*.
- *Alternative measure of redundancy.* Redundancy of review *r* is operationalized as the average number of reviews of all features in review hierarchy of *r* that are co-mentioned in five or more reviews of the same product.

Second, we perform subsample analysis by splitting the data based on product type into search goods and experience goods. The results are reported in Table 6, which show that breadth and depth have positive effects on the review helpfulness regardless of product type, yet the effect of redundancy is only significant for search goods. The insignificance may be explained from the following aspects. There has been evidence for the differing impacts on review helpfulness across product categories [Mudambi and Schuff 2010; Zhu and Zhang 2010]. Experience goods, by nature, require consumers to experience or consume them to truly determine product characteristics [Nelson 1970]. Experience is subjective and thus susceptible to individual differences. Thus, even if different consumers comment on the same features of experience goods, their personal experiences are likely to be different, diminishing the effect of redundancy.

Third, while it is important for our research objective to focus on reviews with more than one vote, selecting reviews based on a cut-off of five or more votes (to increase the chance of receiving a helpful vote) is somewhat arbitrary. We examine the models in a sample of reviews that have three or more votes. The results remain the same in terms of the significance and direction of the relationships except that the effect of depth becomes slightly weaker. In summary, the results of robustness tests provide consistent evidence for the robustness of our findings; additionally, they suggest that the effects of redundancy vary with the product type.

6 DISCUSSION

The results of this study provide strong evidence in support of the framework of semantic hierarchy of product features for review helpfulness assessment. Based on the proposed semantic hierarchy and supporting theories, we introduce three new factors, including breadth, depth, and redundancy, to improve the understanding of review helpfulness. Specifically, breadth and depth are found to have positive effects, and redundancy have a curvilinear effect, on review helpfulness. The robustness test results further confirm the robustness of our model.

6.1 Research Contributions and Implications

This study embodies knowledge moments of multiple genres of inquiry in design science research [Baskerville et al. 2015] and makes multifold research contributions. The semantic hierarchical framework of product features is a form of nomothetic scientific knowledge, and the knowledge is generalizable across different product domains. The framework can serve as a guide in discovering new factors of information quality. Strong information-quality signals are embedded in the semantic hierarchy that go beyond the bag-of-words representation of text. The constructed hierarchies of product features are validated in terms of consistency and reliability.

The proposed research model of review helpfulness, consisting of breadth, depth, and redundancy as main factors, manifests idiographic scientific knowledge that contributes to improved understanding of the helpfulness of online product reviews. The three new factors exemplifiy an application of the semantic hierarchy in assessing the helpfulness of online reviews by investigating the semantic relationships between product features for the first time. The research model is tested and confirmed using regression analysis and robustness checks on online reviews collected from Amazon.com. The transferability of the proposed model is evidenced by the similarities in test results between search goods and experience goods. The comparison also reveals differences, offering a new lens for examining product differentiation in information systems and marketing research. Another moment of idiographic scientific knowledge is a demonstration of practical impacts of the identified factors with the improved performance of predictive models for review helpfulness. A third moment is our novel extension and integration of a multitude of information systems, information science, and marketing theories for the explanation of online review helpfulness. The theoretical extensions also enable reconceptualizing review helpfulness from the viewpoint of prospective consumers. A fourth moment of idiographic scientific knowledge is our expansion of the definition of product features by incorporating service-oriented features (e.g., return policy and product delivery), which have significant presence in online reviews.

This research exhibits multiple moments of idiographic design knowledge. There are the design of approaches to the semantic hierarchy construction, the operationalization of breadth, depth, and redundancy in the context of online reviews, and the performance evaluation of predictive models for review helpfulness (ratio).

6.2 Practical Implications

The findings of this study have practical implications for different stakeholders of online consumer reviews. The proposed model for review helpfulness assessment can be directly used to address the problem of insufficient helpfulness votes. Specifically, by building a predictive model using our proposed features, a website would be able to highlight the most potentially helpful reviews even if they do not yet have or only have a few helpfulness votes. For those reviews, the website would be able to further delineate whether the low vote count is due to review recency or low helpfulness. Furthermore, it enables those websites that do not provide the helpfulness voting function to support consumer purchase decisions by sorting their reviews based on the predicted helpfulness. In addition, the model can aid prospects in screening for helpful online consumer reviews, guide reviewers in preparing helpful reviews, and facilitate online retailers in promoting helpful reviews and even reviewers, which in turn fosters the overall "crowdness" and viability of online review communities. The methods, techniques, and models developed in this study can be used to build consumer decision support tools in online review platforms. Furthermore, the insights gained into product features based on consumer experience can facilitate manufacturers in improving product offerings by enhancing popular product features and correcting defective ones.

Based on the standardized regression coefficients, redundancy is ranked as the top and depth as the fifth most important factors, among other variables included in our research model. Thus, incorporating these factors as input features to build predictive models can potentially enhance the performance in automatic review helpfulness assessment.

7 CONCLUSION AND FUTURE WORK

We introduce and validate three novel factors—review breadth, depth, and redundancy, for assessing the helpfulness of online consumer reviews. The conceptualization of these factors drew on multiple theories and the proposed semantic hierarchy of product features. The findings of this study show that review breadth and depth have positive effects, and redundancy has a curvilinear effect, on review helpfulness. Theoretically, the findings of this study shed light on the significant role of product features in reducing product uncertainty by increasing product diagnosticity for prospective consumers. Practically, this study provides concrete suggestions on how online retailers may leverage variables drawn from the semantic hierarchy of product features to identify and recommend helpful reviews to foster transactional and/or loyalty-building relationships with consumers.

There are several limitations that present opportunities for future research. One is the exclusion of reviewer-related factors and sentiments expressed in online reviews. Incorporating factors such as review reputation and discrete emotions [Ren and Hong 2018] can be promising for review helpfulness assessment. Additional factors can also be derived from the semantic hierarchy of product features to enhance the proposed research model. A second limitation concerns the heterogeneity between different review consumers and online review platforms. Testing the proposed model with a different online review platform and personalizing reviews to preferences of individual consumers are important research questions for the next step. Third, product feature extraction can benefit from recent development in deep learning and word embedding techniques. Fourth, some additional variables such as star rating variance [Danescu-Niculescu-Mizil et al. 2009] and price and interaction effects can be incorporated to improve model fit. Last but not least, based on our data collection, endogeneity unlikely poses a concern to our research model. Yet such a possibility may still exist if an online review is manipulated for helpfulness votes, which calls for future investigations.

REFERENCES

N. Archak, A. Ghose, and P. G. Ipeirotis. 2011. Deriving the pricing power of product features by mining consumer reviews. *Manage. Sci.* 57, 8 (2011) 1485–1509. DOI: 10.1287/mnsc.1110.1370

R. L. Baskerville, M. Kaul, and V. C. Storey. 2015. Genres of inquiry in design-science research: Justification and evaluation of knowledge production. *Manage. Info. Syst. Quart.* 39, 3 (2015), 541–564. DOI: 10.25300/misq/2015/39.3.02

A. Benlian, R. Titah, and T. Hess. 2012. Differential effects of provider recommendations and consumer reviews in e-commerce transactions: An experimental study. *J. Manage. Info. Syst.* 29, 1 (2012), 237–272. DOI: 10.2753/mis0742-1222290107

C. R. Berger and R. J. Calabrese. 1975. Some explorations in initial interaction and beyond: Toward a developmental theory of interpersonal communication. *Hum. Commun. Res.* 1, 2 (1975), 99–112. DOI: 10.1111/j.1468-2958.1975.tb00258.x

E. Borgida and R. E. Nisbett. 1977. The differential impact of abstract vs. concrete information on decisions. *J. Appl. Soc. Psychol.* 7, 3 (1977), 258–271. DOI: 10.1111/j.1559-1816.1977.tb00750.x

K. P. Burnham and D. R. Anderson. 2002. *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach* (2nd ed). Springer-Verlag, 60–65.

C. C. Chen and Y.-D. Tseng. 2011. Quality evaluation of product reviews using an information quality framework. *Decis. Supp. Syst.* 50, 4 (2011), 755–768. DOI: https://doi.org/10.1016/j.dss.2010.08.023

J. A. Chevalier and D. Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *J. Market. Res.* 43, 3 (2006), 345–354. DOI: 10.1509/jmkr.43.3.345

A. Y. K. Chua and S. Banerjee. 2016. Helpfulness of user-generated reviews as a function of review sentiment, product type and information quality. *Comput. Hum. Behav.* 54 (2016), 547–554. DOI: https://doi.org/10.1016/j.chb.2015.08.057

E. K. Clemons, G. G. Gao, and L. M. Hitt. 2006. When online reviews meet hyperdifferentiation: A study of the craft beer industry. *J. Manage. Info. Syst.* 23, 2 (2006), 149–171.

V. P. Crawford and J. Sobel. 1982. Strategic information transmission. *Econometrica* 50, 6 (1982), 1431–1451. DOI: 10.2307/1913390

C. Danescu-Niculescu-Mizil, G. Kossinets, J. Kleinberg, et al. 2009. How opinions are received by online communities: A case study on amazon.com helpfulness votes. In *Proceedings of the 18th International Conference on World Wide Web*. 141–150. DOI: 10.1145/1526709.1526729

A. Dimoka, Y. Hong, and P. A. Pavlou. 2012. On product uncertainty in online markets: Theory and evidence. *Manage. Info. Syst. Quart.* 36, 2 (2012), 395–426.

R. Filieri, F. McLeay, B. Tsui, et al. 2018. Consumer perceptions of information helpfulness and determinants of purchase intention in online consumer reviews of services. *Info. Manage.* (2018). DOI: https://doi.org/10.1016/j.im.2018.04.010

D. A. Garvin. 1984. What does “product quality” really mean? *MIT Sloan Manage. Rev.* 26, 1 (1984).

M. Ghasemaghaei and K. Hassanein. 2015. Online information quality and consumer satisfaction: The moderating roles of contextual factors–A meta-analysis. *Info. Manage.* 52, 8 (2015), 965–981. DOI: https://doi.org/10.1016/j.im.2015.07.001

A. Ghose and P. G. Ipeirotis. 2011. Estimating the helpfulness and economic impact of product reviews: Mining text and reviewer characteristics. *IEEE Trans. Knowl. Data Eng.* 23, 10 (2011), 1498–1512. DOI: 10.1109/TKDE.2010.188

A. R. Hevner, S. T. March, J. Park, et al. 2004. Design science in information systems research. *Manage. Info. Syst. Quart.* 28, 1 (2004), 75–105.

H. Hong, D. Xu, G. A. Wang, et al. 2017. Understanding the determinants of online review helpfulness: A meta-analytic investigation. *Decis. Supp. Syst.* 102 (2017), 1–11. DOI: https://doi.org/10.1016/j.dss.2017.06.007

N. Hu, I. Bose, and N. S. Koh. 2012. Manipulation of online reviews: An analysis of ratings, readability, and sentiments. *Decis. Supp. Syst.* 52, 3 (2012), 674–684. DOI: https://doi.org/10.1016/j.dss.2011.11.002

L. Huang, C.-H. Tan, and W. Ke. 2013. Comprehension and assessment of product reviews: A review-product congruity proposition. *J. Manage. Info. Syst.* 30, 3 (2013), 311–343. DOI: 10.2753/MIS0742-1222300311

M. L. Jensen, J. M. Averbeck, and Z. Zhang. 2013. Credibility of anonymous online product reviews: A language expectancy perspective. *J. Manage. Info. Syst.* 30, 1 (2013), 293–324. DOI: 10.2753/MIS0742-1222300109

N. Jindal and B. Liu. 2007. Analyzing and detecting review spam. In *Proceedings of the 7th IEEE International Conference on Data Mining (ICDM'07)*. 547–552. DOI: 10.1109/ICDM.2007.68

Y. Kang and L. Zhou. Longer is better? A case study of product review helpfulness prediction. In *Proceedings of the 22nd Americas Conference on Information Systems*.

Y. Kang and L. Zhou. 2017. RubE: Rule-based methods for extracting product features from online consumer reviews. *Info. Manage.* 54, 2 (2017), 166–176. DOI: https://doi.org/10.1016/j.im.2016.05.007

Y. Kang, L. Zhou, and D. Zhang. 2014. An integrated method for hierarchy construction of domain-specific terms. In *Proceedings of the IEEE/ACIS 13th International Conference on Computer and Information Science (ICIS'14)*. 485–490. DOI: 10.1109/ICIS.2014.6912181

S. Karimi and F. Wang. 2017. Online review helpfulness: Impact of reviewer profile image. *Decis. Supp. Syst.* 96 (2017), 39–48. DOI: https://doi.org/10.1016/j.dss.2017.02.001

S.-M. Kim, P. Pantel, and T. Chklovski. 2006. Automatically assessing review helpfulness. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 423–430.

K. K. Y. Kuan, K.-L. Hui, and P. Prasarnphanich. 2015. What makes a review voted? An empirical investigation of review voting in online review systems. *J. Assoc. Info. Syst.* 16, 1 (2015).

N. Kumar and I. Benbasat. 2006. Research note: The influence of recommendations and consumer reviews on evaluations of websites. *Info. Syst. Res.* 17, 4 (2006), 425–439. DOI: 10.1287/isre.1060.0107

M. Li, L. Huang, and C.-H. Tan. 2013. Helpfulness of online product reviews as seen by consumers: Source and content features. *International J. Electron. Comm.* 17, 4 (2013), 101–136. DOI: 10.2753/JEC1086-4415170404

X. Li, L. M. Hitt, and Z. J. Zhang. 2011. Product reviews and competition in markets for repeat purchase products. *J. Manage. Info. Syst.* 27, 4 (2011), 9–42. DOI: 10.2753/MIS0742-1222270401

T.-P. Liang, H.-J. Lai, and Y.-C. Ku. 2006. Personalized content recommendation and user satisfaction: Theoretical synthesis and empirical findings. *J. Manage. Info. Syst.* 23, 3 (2006), 45–70.

B. Liu. 2011. *Web Data Mining: Exploring Hyperlinks, Contents, and Usage Data* (2nd ed). Springer-Verlag, Berlin.

X. Lu, S. Ba, and L. Huang. 2013. Promotional marketing or word-of-mouth? Evidence from online restaurant reviews. *Info. Syst. Res.* 24, 3 (2013), 596–612. DOI: 10.1287/isre.1120.0454

X. Lu and X. Zhao. 2014. Differential effects of keyword selection in search engine advertising on direct and indirect sales. *J. Manage. Info. Syst.* 30, 4 (2014), 299–326. DOI: 10.2753/MIS0742-1222300411

M. S. I. Malik and A. Hussain. 2018. An analysis of review content and reviewer variables that contribute to review helpfulness. *Info. Process. Manage.* 54, 1 (2018), 88–104. DOI: https://doi.org/10.1016/j.ipm.2017.09.004

P. M. Markopoulos and E. K. Clemons. 2013. Reducing buyers' uncertainty about taste-related product attributes. *J. Manage. Info. Syst.* 30, 2 (2013), 269–299. DOI: 10.2753/MIS0742-1222300210

J. McAuley and J. Leskovec. 2013. Hidden factors and hidden topics: Understanding rating dimensions with review text. In *Proceedings of the 7th ACM Conference on Recommender Systems*. 165–172. DOI: 10.1145/2507157.2507163

R. K. Merton. 1968. The matthew effect in science. *Science* 159, 3810 (1968), 56.

G. A. Miller. 1956. The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychol. Rev.* 63, 2 (1956), 81–97.

S. M. Mudambi and D. Schuff. 2010. What makes a helpful online review? A study of customer reviews on amazon.com. *Manage. Info. Syst. Quart.* 34, 1 (2010), 185–200.

K. B. Murray. 1991. A test of services marketing theory: Consumer information acquisition activities. *J. Market.* 55, 1 (1991), 10–25. DOI: 10.2307/1252200

P. Nelson. 1970. Information and consumer behavior. *J. Polit. Econ.* 78, 2 (1970), 311–329.

M. P. O'Mahony and B. Smyth. A classification-based review recommender. In *Research and Development in Intelligent Systems*, Vol. 26. 49–62.

P. A. Pavlou, H. Liang, and Y. Xue. 2007. Understanding and mitigating uncertainty in online exchange relationships: A principal-agent perspective. *Manage. Info. Syst. Quart.* 31, 1 (2007), 105–136.

M. I. Posner. 1989. *The Foundations of Cognitive Science.* The MIT Press, Cambridge, MA.

G. Qiu, B. Liu, and J. Bu. 2011. Opinion word expansion and target extraction through double propagation. *Comput. Linguist.* 37, 1 (2011), 9–27. DOI: 10.1162/coli_a_00034

A. R. Rao and K. B. Monroe. 1989. The effect of price, brand name, and store name on buyers' perceptions of product quality: An integrative review. *J. Market. Res.* 26, 3 (1989), 351–357. DOI: 10.2307/3172907

G. Ren and T. Hong. 2018. Examining the relationship between specific negative emotions and the perceived helpfulness of online reviews. *Info. Process. Manage.* (2018). DOI: https://doi.org/10.1016/j.ipm.2018.04.003

A. G. Sawyer, C. Janiszewski, and H. Noel. 2003. A meta-analysis of the spacing effect in verbal learning: Implications for research on advertising repetition and consumer memory. *J. Consum. Res.* 30, 1 (2003), 138–149. DOI: 10.1086/374692

M. Schindler Robert and B. Bickart. 2012. Perceived helpfulness of online consumer reviews: The role of message content and style. *J. Consum. Behav.* 11, 3 (2012), 234–243. DOI: 10.1002/cb.1372

W. Schnotz and C. Kürschner. 2007. A reconsideration of cognitive load theory. *Edu. Psychol. Rev.* 19, 4 (2007), 469–508. DOI: 10.1007/s10648-007-9053-4

R. J. Senter and E. A. Smith. 1967. *Automated Readability Index.* AMRL-TR-6620, Wright-Patterson Air Force Base.

C. E. Shannon. 2001. A mathematical theory of communication. *SIGMOBILE Mob. Comput. Commun. Rev.* 5, 1 (2001), 3–55. DOI: 10.1145/584091.584093

J. P. Singh, S. Irani, and N. P. Rana. 2017. Predicting the "helpfulness" of online consumer reviews. *J. Bus. Res.* 70 (2017), 346–355. DOI: https://doi.org/10.1016/j.jbusres.2016.08.008

R. Ullah, N. Amblee, and W. Kim. 2016. From valence to emotions: Exploring the distribution of emotions in online product reviews. *Decis. Supp. Syst.* 81 (2016), 41–53. DOI: https://doi.org/10.1016/j.dss.2015.10.007

H. R. Unnava and R. E. Burnkrant. 1991. Effects of repeating varied ad executions or brand name memory. *J. Market. Res.* 28, 4 (1991), 406–416. DOI: 10.2307/3172781

J. E. Urbany, P. R. Dickson, and W. L. Wilkie. 1989. Buyer uncertainty and information search. *J. Consum. Res.* 16, 2 (1989), 208–215. DOI: 10.1086/209209

J. G. Walls, G. R. Widmeyer, and O. A. El Sawy. 1992. Building an information system design theory for vigilant EIS. *Info. Syst. Res.* 3, 1 (1992), 36–59. DOI: 10.1287/isre.3.1.36

R. Y. Wang and D. M. Strong. 1996. Beyond accuracy: What data quality means to data consumers. *J. Manage. Info. Syst.* 12, 4 (1996), 5–33. DOI: 10.1080/07421222.1996.11518099

D. Weathers, S. D. Swain, and V. Grover. 2015. Can online product reviews be more helpful? Examining characteristics of information content by product type. *Decis. Supp. Syst.* 79 (2015), 12–23. DOI: http://dx.doi.org/10.1016/j.dss.2015.07.009

C. Wu, H. Che, and T. Y. Chan. 2015. The economic value of online reviews. *Market. Sci.* 34, 5 (2015), 739–754. DOI: 10.1287/mksc.2015.0926

D. Yin, S. D. Bond, and H. Zhang. 2014. Anxious or angry? Effects of discrete emotions on the perceived helpfulness of online reviews. *Manage. Info. Syst. Quart.* 38, 2 (2014), 539–560. DOI: 10.25300/misq/2014/38.2.10

D. Yin, S. Mitra, and H. Zhang. 2016. Research note—When do consumers value positive vs. negative reviews? An empirical investigation of confirmation bias in online word of mouth. *Info. Syst. Res.* 27, 1 (2016), 131–144. DOI: 10.1287/isre.2015.0617

Y. Zheng, K. Zhao, and A. C. Stylianou. 2011. The formation of social influence in online recommendation systems: A study of user reviews on Amazon.com. In *Proceedings of the 32nd International Conference on Information Systems 26*. https://aisel.aisnet.org/icis2011/proceedings/onlinecommunity/26.

L. Zhou and P. Chaovalit. 2008. Ontology-supported polarity mining. *J. Am. Soc. Info. Sci. Technol.* 59, 1 (2008), 98–110. DOI: 10.1002/asi.v59:1

S. Zhou and B. Guo. 2017. The order effect on online review helpfulness: A social influence perspective. *Decis. Supp. Syst.* 93 (2017), 77–87. DOI: https://doi.org/10.1016/j.dss.2016.09.016

F. Zhu and X. Zhang. 2010. Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics. *J. Market.* 74, 2 (2010), 133–148. DOI: 10.1509/jmkg.74.2.133

Received July 2018; revised August 2019; accepted September 2019