The Dark Side of Images: Effect of Customer Generated Images on Product Assessment

Short paper

Yue Guan
School of Economics and Management
Tsinghua University
Beijing, China, 100084
guany.15@sem.tsinghua.edu.cn

Yong Tan
Foster School of Business
University of Washington
Seattle, WA 98195-3226, USA
ytan@uw.edu

Qiang Wei
School of Economics and Management
Tsinghua University
Beijing, China, 100084
weiq@sem.tsinghua.edu.cn

Guoqing Chen
School of Economics and Management
Tsinghua University
Beijing, China, 100084
chengq@sem.tsinghua.edu.cn

Abstract

Customer Generated Image (CGI) on e-commerce platforms has been widely recognized as a marketing tool to persuade customers into purchases. Despite its persuasive power, the effect of CGI on post purchase satisfaction has seldom been examined. This study draws upon Elaboration Likelihood Model and proposes that the affective cues in CGI could distract consumer's cognitive information processing and lead to unsatisfactory purchases with a larger probability. To empirically test our hypothesis, we employed a difference-in-differences model with propensity score weighting method and deep learning based face detection algorithm and found that CGI could cause subsequent review ratings 0.12 stars lower compared with those not exposed to CGI. Additional analysis indicated that this negative effect could be attenuated if the CGI contains human faces or the image review has a low rating. These findings have important implications for online platforms to better leverage user generated rich media content.

Keywords: Customer Generated Image, Elaboration Likelihood Model, affective cues, difference-in-differences, face detection

Introduction

With the surge of social media, live streaming and short video platforms in recent years, numerous images and videos are generated by users every day. According to statistics[1], in one minute's time, YouTube users watch 4,333,560 videos, and Instagram users post 49,380 photos. These user generated images can also act as an important marketing tool to attract new sales. For instance, on Instagram or WeChat, many brands would organize campaigns and encourage their customers to post pictures about their experience with the products. In Amazon Vine Program[2], participating vendors would provide some top-ranking reviewers with free product samples and invite them to post their opinions with images on the platform.

In spite of the rising importance of customer generated images in the retailing industry, to the best of our knowledge, no research has studied the impact of Customer Generated Image (CGI for short) on other customers' product assessment. In academia, much attention has been devoted to investigating the effect

[1] https://www.domo.com/learn/data-never-sleeps-6

[2] https://www.amazon.com/gp/vine/help?ie=UTF8

of product reviews (Hu et al. 2017; Ho et al. 2017; Lee et al. 2015; Yin et al. 2016) or the factors involved in review generation (Huang et al. 2017; Lin et al. 2019), while most of them focus on the textual content and numerical ratings of product reviews with the effectiveness of CGI seldomly examined. There are also some image related research coming up recently with the prosper of deep learning, while they focused on advertisement images (So and Oh, 2018; Zhang et al. 2018) and their persuasive power over customers into purchases (Wang et al. 2016), with little attention paid on the role of CGI and post purchase satisfaction.

It is well recognized that images have perceptual and persuasive advantages over texts (Peracchio and Meyers-Levy 2005). They are generally more expressive, possess high attention grabbing qualities and are remembered better (Childers and Houston 1984). CGI integrates the characteristics of both User Generated Content and visual images, making it more influential than seller-provided information (Goh et al. 2013) and more impressive than textual content generated by customers (Wang et al. 2016). However, does persuasive power equal to post purchase satisfaction? Is more information always a good thing for customers? Located besides review texts, they could act as a distraction to consumer's cognitive information processing related to product quality and fitness (Petty and Cacioppo 1986). Moreover, CGIs are usually posted by consumers without much expertise on taking pictures, thus having a lower aesthetic quality than seller-provided pictures. In other words, CGI could split customers' attention without much value added (Mayer and Moreno 1998).

Based on the research gap above and practical necessity, our research questions in this study are: (1) Does CGI help customers make a more satisfactory purchase decision? (2) How does the effect of CGI differ on different conditions (human faces, review rating, image aesthetic level, etc)? Based on a real world dataset from a large e-commerce platform, we employed a difference-in-differences model with propensity score weighting, and found that CGI actually has a negative effect on subsequent product ratings, meaning that customers are less satisfied with their purchase experience if there's a CGI appearance before purchase. This negative effect is even more prominent after controlling for product fixed effect and time fixed effect. Then we developed the model to a more granular level and examined the effects of specific elements in the image review[3]. Especially, we deployed a state-of-the-art deep learning algorithm to detect whether there are human faces in a CGI. The results showed that this negative effect can be attenuated if the image review has a low rating or there are human faces in the CGI.

Theoretical Background and Hypothesis Development

The prosper of deep learning and artificial intelligence provides us with tools and models to deal with large scale imagery data (LeCun et al. 2015). There are some studies leveraging deep neural networks to conduct visual analytics (Guan et al. 2019) and uncover the relationship between image content and business outcomes. Liu et al. (2018) proposed a model to discover brand attributes from images posted by Instagram users. So et al. (2018) and Wang et al. (2016) demonstrated the effect of product images on customer click behavior and conversion. Zhang et al. (2018) indicated that professional and high quality house images on Airbnb significantly boosted customer demand. Cyr et al. (2009) proved that human images on the website increased image appeal and social presence, reduced ambiguity and risk and consequently enhanced customer trust. Nevertheless, most of them discussed advertisement images generated by the seller side and neglect the role of the numerous user generated images. In a lab experiment, Xu et al. (2015) demonstrated the importance of different product review representation formats and found that video reviews could increase review persuasiveness and purchase intensions. However, the purchase intensions cannot reflect post-purchase product evaluations. Hence it is deemed important to examine the effect of CGIs on product evaluations in a real world e-commerce setting.

According to Elaboration Likelihood Model (ELM) (Petty et al. 1983; Petty and Cacioppo 1986), consumer attitude is formed through two different routes: the central route and peripheral route. In the central route, attitude results from a person's diligent consideration of information that he/she feels is central to the true merits of a particular attitude position. While in the peripheral route, an attitude occurs not because of a thorough consideration of the pros and cons of the issue, but because it is related to positive or negative cues. In the online shopping context, we postulate that both the central and peripheral routes could affect customer attitude, purchase and evaluation decisions.

[3] An image review is defined as a product review containing a Customer Generated Image in this study

A CGI increases the vividness of a product, hence it could quickly grab a consumer's attention before he/she read the textual content of a review (Peracchio and Meyers-Levy 2005). This increased emotional arousal could act as affective cues (Darke et al. 2016) for peripheral-route-based decision making and generate a temporary favorable attitude and purchase intention. In other words, for potential customers, these visual cues act as a credible sales assistant indicating the product's quality thus generate a stronger purchase desire (Xu et al. 2015) than product reviews without a CGI. Yet these affective cues may be irrelevant with the true qualities of a product and the informativeness of CGI remains to be debated (Chen and Xie 2008). As Petty and Cacioppo (1986) shows, the favorable attitude through peripheral route could be temporary and less persistent. At the time of writing a review, the actual product usage experience could lead to the attitude changing toward unfavorable direction.

From an informative perspective, CGI usually has an inferior quality compared with advertisement images as the reviewers are not professional photographers, therefore its informative role remains to be debated. Moreover, as consumer's cognitive resources and time spent on browsing each potential product are limited. With customers relying more on the visual cues of CGI, less attention (Mayer and Moreno 1998) is paid to other product attributes and quality scrutiny, thus customers have a higher probability of making purchase errors. Therefore, we hypothesize that customer satisfaction could be lower with a CGI appearing in the review page before purchase.

H1: Given other conditions the same, product reviews with CGI have a lower subsequent rating compared with product reviews without CGI.

The rating of an image review is highly correlated with the content and quality of a CGI. Take clothes as an example, a low rating implies a disappointment attitude with the product and usually the reviewer would show some defects of the products in a CGI, while a high rating is usually correlated with a CGI in which customers show themselves great fit with the clothes. Therefore, high rating CGIs convey more positive affective cues than low rating CGIs. They send a signal of high quality to uninformed customers and are more prone to persuade customers into hassle purchase. However, these customers did not conduct much cognitive deliberation on whether this product is a fit for them, which consequently lead to more dissatisfaction afterwards than low rating CGIs. Thus we have the following hypothesis.

H2: The effect of CGI on subsequent product ratings is more negative when the image review has a high rating.

One of the most notable differences with CGIs lie in the existence of face identities. Some include face of the reviewer, while others do not. Cyr et al. (2009) demonstrated that a face identity could increase the awareness of the other person in a communication and make the online environment act more like a face-to-face communication. Despite the negative effect of CGI as demonstrated above, we argue that CGI with a face identity could partially alleviate that effect with the following reasons. When there is a face identity in a CGI, the review community acts more like a social network. As Huang et al. (2017) showed that people tend to become more agreeable in a social network with less negative words posted, reviews following a CGI review with face identity will also be less negative than reviews without a face identity. Moreover, images with human faces are less likely to be a fake review, thus increase the credibility of the review content and trust of the subsequent reviewers (Luca and Zervas 2016). Thus we postulate the following hypothesis.

H3: The effect of CGI on subsequent product ratings is less negative when there are human faces in the CGI.

Quantifying the Effect of CGI on Product Ratings

Data Preparation and Key Variables

Our primary empirical analysis utilizes data gathered from Amazon (www.amazon.com). Amazon is one of the largest e-commerce platforms in the world on which tens of thousands of reviews are generated every day. We choose this context because it is representative of the online shopping environment and the data size is large enough to empirically test our research model. Our dataset covers products within a specific category, namely Women casual dresses, which includes all of the product related information and their corresponding product reviews and CGIs. Initially the dataset has 19,149 products and 219,017 reviews, including 15,006 CGIs. As CGIs on the first page and last page could have different probabilities of being noticed (Haugtvedt and Wegener 1994), to alleviate the impact of review order and multiple CGIs, we choose the products with at least 8 reviews and only keep the first 8 reviews for each product in our empirical analysis (on Amazon platform, the first page can only display 8 reviews). Additionally we also delete products with more than one image reviews to exclude the complex effect of multiple CGIs.

In addition to CGI, there are other variables that could potentially affect product ratings. To better deal with the confounding effect (Lee et al. 2015; Moe and Schweidel 2012), three categories of variables are extracted from the dataset to be included in our model, namely, product properties, rating environment variables, and time variables. The summary statistics and corresponding explanations of these variables are presented in Table 1. We also plot the rating distribution for image reviews. Figure 1 shows that 60% of the image reviews concentrate in the 5-star rating group and 80% have at least 4 stars.

Variables	**Variable Explanation**	**Mean**	**Std. Dev.**	**Min**	**Max**
price	The price of a product shown on the website (in US dollars)	32.76	30.18	2.48	306.99
brand	Categorical variable	17.08	5.79	1	20
fitness	Categorical variable (1 represents fit, 2 for large, 3 for small and 4 for unknown)	1.32	0.75	1	4
avg_rating	Product average rating, as a proxy for product quality	3.81	0.54	1.1	5
total_review	Total number of reviews until April 1st, 2017	77.05	168.49	8	3170
product_date	The number of days between product first available date and June 1, 2017	652.56	405.12	65	4328
sum_ review_b$_{it}$	The number of reviews posted for product i before time t	3.44	2.30	0	7
avg_rating_b$_{it}$	The average rating for product i before time t	3.50	1.59	0	5
avg_review_len$_{it}$	The average length of previous reviews for product i before time t	31.40	29.21	0	585
avg_title_len$_{it}$	The average length of review titles for product i before time t	3.81	2.66	0	25
review_date	The number of days between review date and June 1st, 2017	512.13	315.13	63	1749
week_no	The number of weeks between review date and August 2nd, 2010	283.88	45.02	107	348

Table 1. Summary Statistics

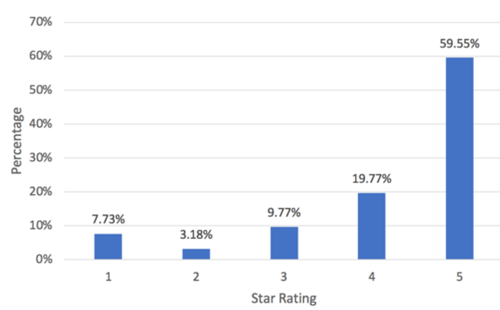

Figure 1. Rating Distribution of Image Reviews

Propensity Score Weighting Method

Following previous literature (Zhang et al. 2017), we employ a difference-in-differences (DID) model to investigate the impact of CGI on subsequent product ratings. All of the products are classified into two groups: the treatment group and control group. A product belongs to the treatment group if there's an image review appearing in the first 8 reviews of a product. After data preprocessing, there are 400 products in the treatment group and 1,863 products in the control group. Considering that the group sizes are unbalanced, and selection problem may exist as the treatment and control group may have different characteristics which could affect our dependent variable of interest, we employ a propensity score weighting method before the DID model. Specifically, a weight (propensity score) is assigned to each data observation to make sure the weighted samples in two groups are balanced on all of the characteristics.

To achieve this, we first run a logistic regression model on all of the products and estimate their propensity scores of being treated accordingly given the product properties. A weight $\mathrm{w_i}$ is derived for each product according to its propensity score ps_i with the formula $\mathrm{w_i} = T/ps_i + (1-T)/ps_i$, where $T = 1$ if product i belongs to the treatment group and $T = 0$ otherwise. Balance check results after reweighting show that the standardized differences are all within 10% level, indicating a negligible sample imbalance (Austin and Stuart 2015).

Main Model Results and Robustness Checks

Our main model is a DID model with weighted least square regression. Unit of analysis is on product-day level. When an image review is posted, the other non-image reviews on the same day are deleted as we do not know their relative arriving order. Our dependent variable is average product rating of product i on time t . The key independent variable is $\mathrm{after\_image_{it}}$ with $\mathrm{after\_image_{it}} = 1$ indicating that there is an image review appearing before time t. treatment = 1 represents the product belonging to the treatment group. In addition, product properties, rating environment and time variables are also controlled in the model. The main model is as shown in Equation (1).

$$Rating_{it} = \beta_0 + \beta_1 \cdot \mathrm{after\_image_{it}} + \beta_2 \cdot \mathrm{treatment_i} + \lambda \cdot ProductProperties_i + \delta \cdot RatingEnvironmentControls_{it} + \gamma \cdot TimeControls_t + \epsilon_{it} \tag{1}$$

The results in Table 2 show that the key variable $\mathrm{after\_image_{it}}$ significantly decreases subsequent rating by 0.117 stars. In other words, the average rating after a CGI in the treatment group is 0.117 stars lower than product ratings in the control group without the presence of a CGI, consistent with our first hypothesis. Despite the small magnitude, this coefficient measures the average effect on every subsequent review. The negative result confirms our theoretical arguments that customers tend to make purchase decisions based on the affective and persuasive cues with the presence of CGIs rather than go through a diligent cognitive information processing. For retailers, although image reviews themselves are usually accompanied with high ratings (as shown in Figure 1), they could be harmful for subsequent product evaluations, which means that this practice of encouraging CGI generation should be used with caution in retail management. The coefficient of *treatment* is significantly positive, demonstrating that products in the treatment group in general have a higher rating than products in the control group.

Regarding product properties, in the group of fitness variables, small size (*fitness=3*) has a significant negative effect on ratings. *Total reviews* also negatively influences product ratings, indicating that a popular product tends to receive more critiques than niche products. Average rating, as a proxy for product quality, has a significantly positive effect on subsequent rating. With respect to the time variable, similar to Li and Hitt (2008), ratings generally follow a downward trend as time goes by (Note that in our setting a larger *review date* implies an earlier review posting date). Among the four rating environment variables, previously posted number of reviews and average title length of previous reviews significantly decrease subsequent ratings, with the former result consistent with Godes and Silva (2012). The negative impact of title length could possibly be due to the differentiation tendency of subsequent reviewers by deliberately giving a low star rating when previous reviews are written elaborately (Lee et al. 2015).

To further justify the results of the main model, a series of alternative models are examined as in Table 2 Model (2)-(4). In model (2), to deal with the concern that relationship between rating and time could be nonlinear, we code the review date into week level and control the week level fixed effect. The results remain similar to the main model with a slight decrease of the $after\_image$ coefficient (-0.114 vs -0.117). Although product properties are controlled in our main model, there may potentially be some variables unobservable that could impact product ratings. Therefore, product fix effects are controlled in model (3). Note that all of the product invariant properties are absorbed in the fixed term. $after\_image$ coefficient is more negative compared with previous models (-0.245 vs -0.117), indicating the main model result as a conservative estimation of CGI's negative impact. In model (4), we control both time and product fix effects. The results are similar to model (3), and the explaining power of model is enhanced greatly, as observed from the outcome of R-square values.

	(1)	(2)	(3)	(4)
Variables	Main Model	Time FE	Product FE	Time & Product FE
after_image	-0.117***	-0.114***	-0.245***	-0.247***
	(0.0293)	(0.0295)	(0.0367)	(0.0373)
treatment	0.0926***	0.0781***		
	(0.0224)	(0.0230)		
avg_rating	0.952***	0.956***		
	(0.0224)	(0.0226)		
fitness-2	-0.0334	-0.0653		
	(0.0704)	(0.0708)		
fitness-3	-0.0986***	-0.0990***		
	(0.0328)	(0.0330)		
fitness-4	-0.00658	-0.00291		
	(0.0975)	(0.0979)		
total_review	-0.000104*	-5.84e-05		
	(6.00e-05)	(6.09e-05)		
price	-0.000421	-0.000377		
	(0.000322)	(0.000327)		
product_ date	-4.34e-05	-4.71e-05		
	(4.29e-05)	(4.40e-05)		
sum_review_b	-0.0142***	-0.0123**	0.0202***	0.0154**
	(0.00503)	(0.00510)	(0.00518)	(0.00642)
avg_rating_b	-0.00225	-0.00527	-0.181***	-0.186***
	(0.00743)	(0.00750)	(0.00918)	(0.00931)
avg_review_len	-0.000106	3.27e-05	0.00322***	0.00342***
	(0.000349)	(0.000353)	(0.000524)	(0.000530)
avg_title_len	-0.0124***	-0.0132***	0.0408***	0.0410***
	(0.00405)	(0.00408)	(0.00598)	(0.00603)
review_date	0.000190***			
	(5.75e-05)			
Constant	0.411***	0.914	4.328***	4.335***
	(0.112)	(1.169)	(0.0238)	(1.261)

Observations	17,430	17,430	17,430	17,430
R-squared	0.173	0.192	0.292	0.311
Week FE	no	yes	yes	yes
Product FE	no	no	yes	yes

Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

Table 2. Main Model Results and Robustness Checks

Heterogeneous Effects of CGI on Product Ratings

This section aims to test H2 and H3 and further understand the conditions of CGI's negative impact on product ratings. Specifically we categorize and discretize image review rating into two groups, high rating group (at least 4 stars) and low rating group (less than 4 stars) and examine if there is any heterogeneity on the outcomes. The results of Model (6) in Table 3 further confirm H2, from which we can observe that the high-rating treatment group has a significant and negative effect while the coefficient of the low-rating treatment group is insignificant, implying that the unsatisfactory evaluations of subsequent customers mainly come from the high-rating image reviews.

To detect human faces from a CGI and test H3, a state-of-the-art face detection algorithm Multi-task Cascaded Convolutional Networks (MTCNN) (Zhang et al. 2016) is deployed to determine the existence of faces and output the corresponding locations in an image, which could achieve validation accuracy of 95.4%. Of the 440 CGIs, faces are detected in 171 images. From the results of Model (7) (Table 3), the negative effect of CGI is reduced by nearly 50% and become insignificant when there are faces in the CGI. This further confirms H3 that the positive emotional arousal by the social presence of others could offset the negative impact of CGIs on product evaluations.

Variables	(5) Baseline	(6) CGI rating	(7) Face
after_image	-0.117***		
	(0.0293)		
after_image_high		-0.135***	
		(0.0310)	
after_image_low		-0.0432	
		(0.0506)	
after_image_face			-0.0617
			(0.0387)
after_image_no_face			-0.151***
			(0.0332)
Constant	0.411***	0.394***	0.428***
	(0.112)	(0.113)	(0.113)
Observations	17,430	17,430	17,430
R-squared	0.173	0.173	0.173

Standard errors in parentheses (*** p<0.01, ** p<0.05, * p<0.1)

Table 3. The Effects of Different CGI Elements on Product Ratings

Conclusions and Future Work

Although extensive research has shown the persuasive advantages of images, this study focus on their effects on post purchase evaluation and shows that CGIs on average lower subsequent product ratings by 0.12 stars. We draw upon Elaboration Likelihood Model and argue that this negative effect is due to the affective cues in CGI which distract consumers' cognitive information processing and decrease decision quality. This explanation is further validated by a more detailed analysis of the heterogeneous treatment effects of CGI. These findings are innovative and enlightening because they are different from previous findings that suggest the positive influences of UGC, such as boosting sales (Chen and Xie 2008), increase trust (Goh et al. 2013), etc. This paper also differs from previous studies demonstrating the positive side of images (Zhang et al. 2018; So et al. 2018; Cyr et al. 2009). Instead, we prove that post purchase satisfaction could deteriorate with the CGI's presence.

Theoretically, this study adds to recent discussion on the potential value and impact of a new kind of user generated content, namely, customer generated images. We focus on the effect of the visual element in the product review and enhance the understanding of visual content's complex marketing power. Moreover, this study draws theoretical support from ELM, explains the effect of CGIs on subsequent ratings and provides empirical evidence on its applicability to this study's context. Finally, this study is the first to demonstrate the complex effects of CGIs in online shopping and review community.

Practically, customer image is a double-edged sword. CGIs are more likely to persuade customers into purchases, but also more likely to incur unsatisfactory purchase experiences. Platform or online retailers should be careful in adopting customer images as a marketing tool, as they may cause unexpected decline of product ratings, which damages the relationship with customers and brand loyalty. To alleviate this negative effect, some measures could be leveraged to encourage an objective and elaborate product evaluation and also encourage more image reviews with face identities.

As this is an ongoing research, for future work, we will dig deeper into the image content, and investigate the influence of CGI under different CGI quantity, quality, and aesthetic levels to better understand the underlying mechanism. Moreover, more robustness checks and alternative econometric models are necessary to further identify and validate the current results. If possible, a lab experiment could also be designed to discover consumers' underlying decision process so as to further validate the theoretical arguments proposed in this study. In this study , we focus on the US market and a specific shopping context, which is fashion, and conduct the empirical analysis. Whether our conclusions could be applied in other contexts, culture or other product categories is another interesting topic worth further investigation.

Acknowledgement

The work was partly supported by the National Natural Science Foundation of China (71772101/71490724) and the MOE Project of Key Research Institute of Humanities and Social Sciences at Universities (17JJD630006).

References

Austin, P. C., and Stuart, E. A. 2015. "Moving towards Best Practice When Using Inverse Probability of Treatment Weighting (IPTW) Using the Propensity Score to Estimate Causal Treatment Effects in Observational Studies," *Statistics in Medicine* (34:28), pp. 3661-3679.

Chen, Y., and Xie, J. 2008. "Online Consumer Review: Word-of-Mouth as a New Element of Marketing Communication Mix," *Management Science* (54:3), pp. 477-491.

Childers, T. L., and Houston, M. J. 1984. "Conditions for a Picture-Superiority Effect on Consumer Memory," *Journal of Consumer Research* (11:2), pp. 643-654.

Cyr, D., Head, M., Larios, H., and Pan, B. 2009. "Exploring Human Images in Website Design: a Multi-Method Approach," *MIS Quarterly* (33:3), pp. 539-566.

Darke, P. R., Chattopadhyay, A., and Ashworth, L. 2006. The Importance and Functional Significance of Affective Cues in Consumer Choice. *Journal of Consumer Research* (33:3), pp. 322-328.

Godes, D., and Silva C.J. 2012. "Sequential and Temporal Dynamics of Online Opinion." *Marketing Science* (31:3), pp. 448-473.
Goh, K., Heng, C., and Lin, Z. 2013. "Social Media Brand Community and Consumer Behavior: Quantifying the Relative Impact of User-and Marketer-Generated Content," *Information Systems Research* (24:1), PP. 88-107.
Guan, Y., Wei, Q., and Chen, G. 2019. "Deep Learning Based Personalized Recommendation with Multi-view Information Integration," *Decision Support Systems* (118), pp. 58-69.
Haugtvedt, C. P., and Wegener, D. T. 1994. "Message Order Effects in Persuasion: An Attitude Strength Perspective," *Journal of Consumer Research* (21:1), pp. 205-218.
Ho, Y. C., Wu, J., and Tan, Y. 2017. "Disconfirmation Effect on Online Rating Behavior: A Structural Model," *Information Systems Research* (28:3), pp. 626-642.
Hu, N., Pavlou, P. A., and Zhang, J. 2017. "On Self-Selection Biases in Online Product Reviews," *MIS Quarterly* (41:2), pp. 449-471.
Huang, N., Hong, Y., and Burtch, G. 2017. "Social Network Integration and User Content Generation: Evidence from Natural Experiments," *MIS Quarterly* (41:4), pp. 1035-1058.
LeCun, Y.A., Bengio, Y., and Hinton, G.E. 2015 , "Deep Learning", *Nature* (521), pp. 436–444.
Lee, Y. J., Hosanagar, K., and Tan, Y. 2015. "Do I Follow My Friends or the Crowd? Information Cascades in Online Movie Ratings," *Management Science* (61:9), pp. 2241-2258.
Li, X., and Hitt, L.M. 2008. "Self-selection and Information Role of Online Product Reviews," *Information Systems Research* (19:4), pp. 456-474.
Lin, Z., Zhang, Y., and Tan, Y. 2019. "An Empirical Study of Free Product Sampling and Rating Bias," *Information Systems Research* (30:1), pp. 260-275.
Liu, L., Dzyabura, D., and Mizik, N. 2018. "Visual Listening In: Extracting Brand Image Portrayed on Social Media," in *Workshops at the Thirty-Second AAAI Conference on Artificial Intelligence,* New Orleans, LA, pp. 71-77.
Luca, M., and Zervas, G. 2016. Fake it Till You Make It: Reputation, Competition, and Yelp Review Fraud. *Management Science*, (62:12), pp. 3412-3427.
Mayer, R. E., and Moreno, R. 1998. "A Split-attention Effect in Multimedia Learning: Evidence for Dual Processing Systems in Working Memory," *Journal of Educational Psychology* (*90*:2), pp. 312-320.
Moe, W. W., and Schweidel, D. A. 2012. "Online Product Opinions: Incidence, Evaluation, and Evolution," *Marketing Science* (31:3), pp. 372-386.
Peracchio, L. A. Meyers-Levy, J. 2005. "Using Stylistic Properties of Ad Pictures to Communicate with Consumers," *Journal of Consumer Research* (32:1), pp. 29–40.
Petty, R. E., Cacioppo, J. T., and Schumann, D. 1983. "Central and Peripheral Routes to Advertising Effectiveness: The Moderating Role of Involvement," J*ournal of Consumer Research* (10:2), pp. 135-146.
Petty, R. E., and Cacioppo, J. T. 1986. "The Elaboration Likelihood Model of Persuasion," in *Communication and Persuasion*, Springer, New York, pp. 1-24.
Shin, D., He, S., Lee, G. M., Whinston, A. B., Cetintas, S., and Lee, K. C. 2016. "Content Complexity, Similarity, and Consistency in Social Media: A Deep Learning Approach," working paper.
So, H., and Oh, W. 2018. "Picture Perfect: An Image Mining of Advertising Content and Its Effects on Social Targeting," in *Proceedings of the 39th International Conference on Information Systems*, San Francisco, CA.
Wang, M., Li, X., and Chau, P.Y. 2016. "The Impact of Photo Aesthetics on Online Consumer Shopping Behavior: An Image Processing-Enabled Empirical Study," in *Proceedings of the 37th International Conference on Information Systems,* Dublin, Ireland.
Xu, P., Chen, L., and Santhanam R. 2015. "Will Video Be the Next Generation of E-commerce Product Reviews? Presentation format and the Role of Product Type," *Decision Support Systems* (73), pp. 85-96.
Yin, D., Mitra, S., and Zhang, H. 2016. Research note—When Do Consumers Value Positive vs. Negative Reviews? An Empirical Investigation of Confirmation Bias in Online Word of Mouth. *Information Systems Research* (27:1), pp. 131-144.
Zhang, S., Lee, D., Singh, P. V., and Srinivasan, K. 2017. "How Much Is an Image Worth? Airbnb Property Demand Estimation Leveraging Large Scale Image Analytics," working paper.
Zhang, K., Zhang, Z., Li, Z., and Qiao, Y. 2016. "Joint Face Detection and Alignment Using Multitask Cascaded Convolutional Networks," *IEEE Signal Processing Letters* (23:10), pp. 1499-1503.