

Original Research Article

The Dark Side of Generative AI in Marketing: AI-Generated False and Misleading Advertising and Its Consumer Consequences

Dr. Pankaj Nandurkar,¹ Mr. Pradeep Gupta,² Mr. Sangram Andhe,³ and Dr. Krishnakant Lasune⁴

¹Professor, MBA, Thakur College of Engineering and Technology, Kandivali East, Mumbai, India

²Assistant Professor, MBA, Thakur College of Engineering and Technology, Kandivali East, Mumbai, India

³Assistant Professor, MBA, Thakur College of Engineering and Technology, Kandivali East, Mumbai, India

⁴Assistant Professor, PGDM, Thakur Global Business School, Kandivali East, Mumbai, India

Received: 10 September, 2026 | Accepted: 24 September, 2026 | Published: 15 October, 2026

Abstract:

Generative artificial intelligence (GenAI) is transforming advertising by enabling marketers to create persuasive text, images, and other promotional content rapidly and at lower cost. At the same time, the ability of generative systems to produce highly realistic content creates concerns about deception, misinformation, transparency, and the credibility of consumer-facing information. The present study examines the consumer consequences of AI-generated false and misleading advertising, with particular emphasis on hallucinated product claims. The study focuses on whether AI-generated hallucinated product claims influence perceived deception, AI skepticism, and perceived information reliability, and whether these consumer responses subsequently shape purchase intention. Drawing on the stimulus-organism-response (S-O-R) framework, the study conceptualizes hallucinated product claims as the stimulus, the three consumer evaluations as the organism, and purchase intention as the response. A quantitative, cross-sectional, deductive research design was adopted. Data were collected from 320 business-school students and faculty members through a structured questionnaire using a five-point Likert scale, and the model was

assessed using SmartPLS. The reported measurement results show Cronbach's alpha values above 0.70, composite reliability above the stated threshold, and average variance extracted above 0.50 for all constructs, indicating acceptable reliability and convergent validity. The reported discriminant-validity matrix also supports construct distinctiveness. The structural model reports a VIF range of 1.47-2.24 and an R^2 of 0.421 for purchase intention. Four of the six hypotheses are supported. AI-generated hallucinated product claims significantly increase perceived deception and significantly reduce perceived information reliability. Perceived deception significantly reduces purchase intention, while perceived information reliability significantly increases purchase intention. The paths from hallucinated product claims to AI skepticism and from AI skepticism to purchase intention are not statistically significant. The findings indicate that message-level judgments of deception and information reliability are more consequential for purchase intention in the reported model than generalized skepticism toward AI. The study extends the S-O-R perspective to the dark side of GenAI in advertising and highlights the importance of factual verification, transparency, human oversight, disclosure, and consumer AI literacy.

Keywords: Generative artificial intelligence; AI hallucinations; misleading advertising; perceived deception; AI skepticism; perceived information reliability; purchase intention; stimulus-organism-response theory.

Introduction

Generative AI is increasingly embedded in marketing communication because it allows marketers to produce persuasive text, images, and other promotional content rapidly and at comparatively low cost. This transformation expands the role of AI from supporting analytical and targeting activities to directly generating consumer-facing advertising content. Grigsby, Michelsen, and Zamudio (2025) examine service advertising in the era of generative AI and emphasize issues surrounding disclosure, trust, and advertising communication, while Feng and Kim (2025) examine predictors of consumer trust in AI-generated personalized advertising. These developments demonstrate that the efficiency and creative potential of GenAI are closely linked to questions of consumer evaluation and trust.

The same generative capability that increases marketing efficiency can also generate false, fabricated, exaggerated, or unsupported information. When such information enters advertising, an AI system may produce product claims that appear credible even though they lack factual support. The problem is especially important in advertising because promotional communication is intended to shape perceptions and influence purchase decisions. A hallucinated product specification, unsupported efficacy claim, or exaggerated product benefit may therefore operate not merely as an information error but as a persuasive message that changes how consumers evaluate the product and the advertiser.

The realism of AI-generated content further complicates the problem. Consumers may struggle to distinguish AI-generated advertising elements from human-created content, creating uncertainty around authenticity, transparency, and credibility. Koning and Voorveld (2025) show that disclosure of AI involvement can alter consumer responses to advertisements and organizations, while the literature summarized in the study indicates that perceived authenticity and moral reactions also influence responses to AI-authored marketing communication. These findings suggest that the effects of AI in advertising are context-dependent and may vary according to how AI is used, whether AI involvement is disclosed, what type of content is produced, the credibility signals surrounding the advertisement, and consumer characteristics.

The present research focuses on a more specific problem than AI-generated advertising in general. It examines advertising that contains AI-generated hallucinated product claims, defined in the study as false, fabricated, exaggerated, or unsupported product information generated by AI. This distinction is important because consumers may accept the use of AI for marketing production while simultaneously rejecting advertising information that they perceive as deceptive or unreliable. The central research issue is therefore the consumer consequence of inaccurate AI-mediated advertising information rather than AI adoption itself.

The study addresses this issue through an S-O-R framework. AI-generated hallucinated product claims form the stimulus, perceived deception, AI skepticism, and perceived information reliability represent the consumer's internal responses, and purchase intention represents the behavioral response. The study proposes that hallucinated claims will increase perceived deception and AI skepticism while reducing perceived information reliability, and that perceived deception and AI skepticism will reduce purchase intention whereas information reliability will increase it.

The study is motivated by a gap identified in the emerging literature. Although hallucinations, misinformation, and unreliable AI outputs are increasingly recognized as epistemic risks in knowledge production, organizational decision-making, and governance, comparatively limited attention has been given to the consumer-facing marketing context. The study therefore seeks to understand what happens when advertising generated or assisted by AI contains false or misleading product information and how consumers respond in terms of deception perceptions, skepticism, information reliability, and purchase intention.

Literature Review

The literature on GenAI in advertising indicates that artificial intelligence has become an important mechanism for generating persuasive and personalized communication. Grigsby, Michelsen, and Zamudio (2025) examine service advertising in the era of generative AI and focus attention on disclosure, trust, and the nature of advertising communication. Feng and Kim (2025) similarly examine consumer trust in AI-generated personalized advertising. Together, these studies suggest that the growing use of AI in advertising creates a dual outcome: marketers obtain greater speed, scalability, and creative flexibility, while consumers must assess whether the resulting communication is trustworthy and credible.

Trust and credibility are central to the evaluation of AI-generated advertising. When consumers encounter content produced by an algorithmic system, they may evaluate both the advertising message and the technology behind the message. Koning and Voorveld (2025) demonstrate that AI disclosure can alter trust in advertisements and organizations, indicating that disclosure does not automatically produce more favorable responses. Instead, awareness of AI involvement may activate persuasion knowledge and change perceptions of the communication. This supports the need to examine consumer judgments at multiple levels rather than treating AI acceptance as a single outcome.

Perceived authenticity also plays an important role in responses to AI-authored marketing communication. Kirk and Givi (2025), as cited in the supplied presentation, focus on the AI-authorship effect and connect AI-generated marketing communication with authenticity, moral disgust, and consumer responses. The relevance to misleading advertising lies in the possibility that inaccurate AI-generated claims may make an advertisement appear not only unreliable but also dishonest or manipulative. Perceived deception therefore provides a mechanism through which false advertising information can reduce favorable consumer responses.

The concept of AI hallucination has received substantial attention as an epistemic risk because generative systems can generate information that sounds plausible but is unsupported or inaccurate. In the present marketing context, the concern is that such information may be embedded in product advertising and used to influence consumer decisions. Unlike a general factual error, a hallucinated advertising claim can become part of a persuasive message, which increases the possibility that consumers will rely on the information when forming product judgments.

Perceived information reliability is consequently important because consumers need dependable information when evaluating products. If an advertisement contains hallucinated claims, the consumer may question the accuracy, trustworthiness, and dependability of the information provided. The present study therefore separates information reliability from perceived deception. A consumer may perceive a particular message as deceptive because of the misleading claim, while simultaneously evaluating the factual content itself as unreliable. Distinguishing these concepts enables a more detailed understanding of the consumer consequences of AI-mediated misinformation.

Purchase intention is treated in the study as the ultimate response variable. The proposed logic is that purchase intention will depend on how consumers interpret the advertising information. When the advertisement is perceived as deceptive, willingness to purchase should decline. When the information is perceived as reliable, willingness to purchase should increase. AI skepticism is also proposed as a negative predictor of purchase intention because greater doubt regarding AI-generated information may discourage consumers from relying on the message. The resulting model connects the information problem created by hallucinated claims with distinct psychological evaluations and a behavioral intention.

The literature therefore leaves an important research gap. Existing studies provide evidence on AI disclosure, trust, authenticity, and general consumer responses to AI-generated advertising, while broader GenAI research recognizes hallucination and misinformation risks. However, limited empirical attention has been directed toward how consumers respond when AI-generated advertising specifically contains false or misleading product claims. The present research addresses this gap by examining three separate consumer responses and their relationships with purchase intention within an S-O-R framework.


Methodology

The study adopts a quantitative, cross-sectional, and deductive research design. The quantitative approach is appropriate for empirically testing the hypothesized relationships among latent constructs, while the cross-sectional design captures respondents' evaluations at a single point in time. The deductive orientation reflects the use of the S-O-R framework to derive and test the proposed relationships between AI-generated hallucinated product claims, consumer responses, and purchase intention.

The study population consists of business-school students and faculty members, and the reported sample contains 320 respondents selected through convenience sampling. The unit of analysis is the individual respondent. Because the presentation does not provide a detailed demographic profile, no additional demographic characteristics are introduced in the present paper. Data were collected through a structured questionnaire using a five-point Likert scale.

The study operationalizes AI-generated hallucinated product claims as false, fabricated, exaggerated, or unsupported product information generated by AI. Perceived deception refers to the extent to which consumers

perceive an advertisement as misleading, dishonest, or deceptive. AI skepticism represents the degree to which consumers doubt or distrust AI-generated information and advertising. Perceived information reliability represents the extent to which consumers perceive the advertising information as accurate, trustworthy, and dependable. Purchase intention refers to the consumer's willingness or likelihood to purchase the advertised product. AI literacy was also included as a measured construct in the measurement assessment, although the supplied presentation does not specify a structural hypothesis involving AI literacy.



The conceptual model places AI-generated hallucinated product claims as the stimulus, perceived deception, AI skepticism, and perceived information reliability as internal consumer responses, and purchase intention as the final response. The model proposes six hypotheses. H1: Generative AI-generated hallucinated product claims have a significant positive effect on perceived deception.

H2: Generative AI-generated hallucinated product claims have a significant positive effect on AI skepticism.

H3: Generative AI-generated hallucinated product claims have a significant negative effect on perceived information reliability.

H4: Perceived deception has a significant negative effect on purchase intention.

H5: AI skepticism has a significant negative effect on purchase intention.

H6: Perceived information reliability has a significant positive effect on purchase intention.

H1 predicts a significant positive effect of AI-generated hallucinated product claims on perceived deception; H2 predicts a significant positive effect on AI skepticism; H3 predicts a significant negative effect on perceived information reliability; H4 predicts a significant negative effect of perceived deception on purchase intention; H5 predicts a significant negative effect of AI skepticism on purchase intention; and H6 predicts a significant positive effect of perceived information reliability on purchase intention.

SmartPLS was used to assess the measurement and structural models. Reliability was evaluated through Cronbach's alpha and composite reliability, convergent validity through average variance extracted (AVE), and

discriminant validity through the reported construct matrix. Structural assessment used path coefficients, t-statistics, p-values, R-squared, and variance inflation factor (VIF). The study uses a Cronbach's alpha threshold above 0.70 and an AVE threshold above 0.50 as the reported acceptance criteria.

Data Analysis

Construct	Cronbach's Alpha	CR (rho_A)	CR (rho_C)	AVE
AI-Generated Hallucinated Product Claims	0.756	0.764	0.838	0.565
Perceived Deception	0.774	0.781	0.852	0.591
AI Skepticism	0.731	0.746	0.821	0.538
Perceived Information Reliability	0.749	0.758	0.833	0.556
Purchase Intention	0.781	0.789	0.858	0.602
AI Literacy	0.738	0.751	0.827	0.548
<i>Acceptable Threshold</i>	<i>> 0.700</i>	<i>> 0.700</i>	<i>> 0.700</i>	<i>> 0.500</i>

The measurement model demonstrates acceptable internal consistency across the constructs. Cronbach's alpha values range from 0.731 to 0.781, which are above the stated threshold of 0.70. For AI-generated hallucinated product claims, Cronbach's alpha is 0.756, composite reliability values are 0.764 and 0.838, and AVE is 0.565. For perceived deception, the corresponding values are 0.774, 0.781, 0.852, and 0.591. AI skepticism records values of 0.731, 0.746, 0.821, and 0.538. Perceived information reliability records 0.749, 0.758, 0.833, and 0.556. Purchase intention records 0.781, 0.789, 0.858, and 0.602, while AI literacy records 0.738, 0.751, 0.827, and 0.548. Since all AVE values exceed 0.50 and all reported reliability values exceed the stated threshold, convergent validity and internal consistency are established according to the assessment reported in the presentation.

Construct	AI-HPC	Perceived Deception	AI Skepticism	Information Reliability	Purchase Intention	AI Literacy
AI-Generated Hallucinated Product Claims	0.752					
Perceived Deception	0.518	0.769				
AI Skepticism	0.472	0.557	0.734			
Perceived Information Reliability	0.441	0.493	0.516	0.746		
Purchase Intention	0.386	0.421	0.397	0.583	0.776	
AI Literacy	0.329	0.351	0.382	0.428	0.365	0.740
Construct	AI-HPC	Perceived	AI	Information	Purchase	AI Literacy

		Deception	Skepticism	n Reliability	Intention	
--	--	------------------	-------------------	----------------------	------------------	--

The reported discriminant-validity matrix also supports construct distinctiveness. The diagonal square-root-of-AVE values are 0.752 for AI-generated hallucinated product claims, 0.769 for perceived deception, 0.734 for AI skepticism, 0.746 for perceived information reliability, 0.776 for purchase intention, and 0.740 for AI literacy. The off-diagonal correlations are lower than the corresponding diagonal values in the reported matrix. For example, AI-generated hallucinated product claims correlate with perceived deception at 0.518, with AI skepticism at 0.472, with information reliability at 0.441, with purchase intention at 0.386, and with AI literacy at 0.329. The presentation therefore concludes that discriminant validity is established for the latent variables.

#	β	T-Stat	P	Decision
H1	0.412	4.247	0.000	Supported
H2	0.148	1.626	0.104	Not Supported
H3	-0.286	2.724	0.007	Supported
H4	-0.231	2.538	0.011	Supported
H5	-0.096	1.103	0.270	Not Supported
H6	0.351	3.545	0.000	Supported

The structural model reports a VIF range of 1.47-2.24, which is below the stated threshold of 5.00 and indicates no material multicollinearity problem. The model reports an R-squared value of 0.421 for purchase intention, meaning that the predictors included in the structural model explain 42.1 percent of the variance in purchase intention in the reported sample. Four of the six proposed hypotheses are statistically supported.

H1 is supported because AI-generated hallucinated product claims have a significant positive effect on perceived deception, with a path coefficient of 0.412, a t-statistic of 4.247, and $p = 0.000$. This indicates that the presence of false, fabricated, exaggerated, or unsupported product claims increases the extent to which consumers perceive the advertisement as deceptive. H2 is not supported because the effect of hallucinated product claims on AI skepticism is positive but statistically insignificant, with $\beta = 0.148$, $t = 1.626$, and $p = 0.104$. The result suggests that misleading claims do not necessarily translate into a broader skeptical orientation toward AI itself.

H3 is supported because AI-generated hallucinated product claims significantly reduce perceived information reliability, with $\beta = -0.286$, $t = 2.724$, and $p = 0.007$. This result indicates that inaccurate AI-generated claims undermine the degree to which consumers view the advertising information as accurate, trustworthy, and dependable. H4 is also supported because perceived deception has a significant negative effect on purchase intention, with $\beta = -0.231$, $t = 2.538$, and $p = 0.011$. The implication is that when consumers perceive the advertisement as deceptive, their willingness or likelihood of purchasing the advertised product declines.

H5 is not supported because AI skepticism does not have a statistically significant effect on purchase intention, with $\beta = -0.096$, $t = 1.103$, and $p = 0.270$. Although the direction is negative, the statistical evidence does not support the proposed relationship. H6 is supported because perceived information reliability has a significant positive effect on purchase intention, with $\beta = 0.351$, $t = 3.545$, and $p = 0.000$. The combined results indicate that consumer responses are more strongly associated with message-level deception and information reliability than with generalized AI skepticism.

Taken together, the data support the central proposition of the study that hallucinated product claims create adverse consumer consequences, but the mechanism is selective. Hallucinated claims clearly increase perceived deception and reduce perceived information reliability, and these two responses have significant relationships with purchase intention in opposite directions. By contrast, AI skepticism is not significantly influenced by the hallucinated claims and is not significantly related to purchase intention. The findings therefore suggest that consumers can reject or distrust the specific advertising message without necessarily developing generalized skepticism toward AI.

Conclusion

This study examined the dark side of generative AI in marketing by focusing on AI-generated false and misleading advertising and, more specifically, hallucinated product claims. The research used an S-O-R framework to explain how misleading AI-generated claims function as an external stimulus that shapes internal consumer evaluations and ultimately influences purchase intention. The measurement assessment reported in the study demonstrates acceptable reliability, convergent validity, and discriminant validity, while the structural model explains 42.1 percent of the variance in purchase intention.

The main findings demonstrate that AI-generated hallucinated product claims significantly increase perceived deception and significantly reduce perceived information reliability. In turn, perceived deception significantly reduces purchase intention, whereas perceived information reliability significantly increases purchase intention. The proposed paths involving AI skepticism are not statistically significant. These results indicate that the consumer consequences of misleading GenAI advertising are more closely tied to perceived deception and the reliability of the information than to a generalized distrust of AI.

The study consequently extends the application of the S-O-R perspective to AI-mediated misleading advertising and contributes to the emerging discussion of the dark side of GenAI in marketing. It demonstrates that the central risk may not be that consumers universally reject AI-generated advertising, but that inaccurate claims undermine the credibility and reliability of the specific marketing information used in purchase evaluation. As GenAI becomes increasingly embedded in advertising processes, accuracy and reliability therefore become central conditions for sustaining consumer confidence.

Practical Implications

The findings have direct implications for marketers. Organizations using GenAI for advertising should strengthen fact-checking and verification procedures, particularly for product specifications, product performance statements, efficacy claims, and other factual assertions. The significant negative relationship between perceived deception and purchase intention indicates that inaccurate claims can have immediate commercial consequences in addition to ethical implications.

Brands should prioritize accuracy and transparency when using AI to create consumer-facing advertising. Agencies should establish human review and validation procedures before AI-generated content is published, especially when the content includes factual or performance-related claims. Human oversight can reduce the likelihood that plausible but unsupported outputs are presented to consumers as verified facts.

Advertising and digital platforms can strengthen AI-content disclosure and verification mechanisms so that consumers receive clearer signals about how advertising content was produced and whether important claims have been checked. Such mechanisms should complement, rather than replace, substantive verification of factual claims.

Disclosure is useful, but the results of the present study emphasize that the reliability of the information itself is central to purchase intention.

The results also support greater attention to consumer AI literacy. Consumers who understand the potential for generative systems to produce inaccurate or unsupported information may be better positioned to critically evaluate advertising claims and seek corroborating evidence before making purchase decisions. Consumer education can therefore form an important part of a broader response to the risks of AI-mediated marketing misinformation.

Limitation

The cross-sectional design limits the ability to establish causal relationships over time. The observed relationships should therefore be interpreted as associations within the reported sample and design rather than definitive evidence of long-term causal effects. Future research can use experimental and longitudinal designs to test how repeated exposure to misleading AI advertising affects consumer perceptions and behavior.

The study relies on self-reported responses, which may be influenced by perception and response biases. Future work could complement survey-based measures with behavioral indicators such as advertisement clicks, product choice, information-search behavior, or actual purchase behavior. The study also examines purchase intention rather than actual purchasing behavior, so the translation from intention to behavior remains to be established.

The research focuses specifically on hallucinated product claims and does not examine other forms of AI-enabled deception, such as deepfake endorsements, synthetic testimonials, manipulated product demonstrations, or fabricated social proof. These forms may activate different psychological mechanisms and should be examined separately in future studies. In addition, brand trust, advertising involvement, prior AI experience, and other potentially relevant factors are not included in the reported structural model.

Finally, the sample consists of business-school students and faculty members selected through convenience sampling. The findings therefore should not automatically be generalized to the wider consumer population. Future research can use more heterogeneous and probability-based samples to assess whether the observed relationships remain stable across consumer groups, product categories, and levels of familiarity with AI.

Bibliography

- Appelman, A., & Sundar, S. S. (2016). Measuring message credibility: Construction and validation of an exclusive scale. *Journalism & Mass Communication Quarterly*, 93(1), 59–79.
<https://doi.org/10.1177/1077699015606057>
- Farquhar, S., Kossen, J., Kuhn, L., & Gal, Y. (2024). Detecting hallucinations in large language models using semantic entropy. *Nature*, 630, 625–630. <https://doi.org/10.1038/s41586-024-07421-0>
- Feng, Y., & Kim, H. J. (2025). Decoding the trust matrix: Unraveling key predictors of consumer trust in AI-generated personalized advertising. *Journal of Interactive Advertising*, 25(2).
<https://doi.org/10.1080/15252019.2025.2468286>
- Friestad, M., & Wright, P. (1994). The persuasion knowledge model: How people cope with persuasion attempts. *Journal of Consumer Research*, 21(1), 1–31. <https://doi.org/10.1086/209380>

- Grewal, D., Saturnino, C. B., Davenport, T., & Guha, A., et al. (2025). How generative AI is shaping the future of marketing. *Journal of the Academy of Marketing Science*, 53, 702–722. <https://doi.org/10.1007/s11747-024-01064-3>
- Grigsby, J. L., Michelsen, M., & Zamudio, C. (2025). Service ads in the era of generative AI: Disclosures, trust, and intangibility. *Journal of Retailing and Consumer Services*, 84, 104231. <https://doi.org/10.1016/j.jretconser.2025.104231>
- Kirk, C. P., & Givi, J. (2025). The AI-authorship effect: Understanding authenticity, moral disgust, and consumer responses to AI-generated marketing communications. *Journal of Business Research*, 186, 114984. <https://doi.org/10.1016/j.jbusres.2024.114984>
- Koning, B., & Voorveld, H. A. M. (2025). Disclaimer! This content is AI-generated: How AI-disclosures influence trust in advertisements and organizations. *Journal of Interactive Advertising*. <https://doi.org/10.1080/15252019.2025.2554149>
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103. <https://doi.org/10.1016/j.obhdp.2018.12.005>
- MacKenzie, S. B., Lutz, R. J., & Belch, G. E. (1986). The role of attitude toward the ad as a mediator of advertising effectiveness: A test of competing explanations. *Journal of Marketing Research*, 23(2), 130–143. <https://doi.org/10.1177/002224378602300205>
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20(3), 709–734. <https://doi.org/10.5465/amr.1995.9508080335>
- Mehrabian, A., & Russell, J. A. (1974). *An approach to environmental psychology*. MIT Press.
- Morhart, F. M., Malär, L., Guèvremont, A., Girardin, F., & Grohmann, B. (2015). Brand authenticity: An integrative framework and measurement scale. *Journal of Consumer Psychology*, 25(2), 200–218. <https://doi.org/10.1016/j.jcps.2014.11.006>
- Ohanian, R. (1990). Construction and validation of a scale to measure celebrity endorsers' perceived expertise, trustworthiness, and attractiveness. *Journal of Advertising*, 19(3), 39–52. <https://doi.org/10.1080/00913367.1990.10673191>
- Tsung-Yu Hou, Y. Y., & Jung, M. F. (2021). Who is the expert? Reconciling algorithm aversion and algorithm appreciation in AI-supported decision making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2), Article 477. <https://doi.org/10.1145/3479864>
- Visentin, M., Pizzi, G., & Pichierri, M. (2019). Fake news, real problems for brands: The impact of content truthfulness and source credibility on consumers' behavioral intentions toward the advertised brands. *International Journal of Advertising*, 38(3), 390–414. <https://doi.org/10.1080/02650487.2018.1535221>
- Weismueller, J., Harrigan, P., Wang, S., & Soutar, G. N. (2020). Influencer endorsements: How advertising disclosure and source credibility affect consumer purchase intention on social media. *Australasian Marketing Journal*, 28(4), 160–170. <https://doi.org/10.1016/j.ausmj.2020.03.002>

- Wu, L., Dodoob, N. A., & Wen, T. J. (2024/2025). Disclosing AI's involvement in advertising to consumers: A task-dependent perspective. *Journal of Advertising*, 54(1), 20–38. <https://doi.org/10.1080/00913367.2024.2309929>
- Xie, Z., Yu, Y., & Chen, M. (2022). The searching artificial intelligence: Consumers show less aversion to algorithm-recommended search products. *Psychology & Marketing*. <https://doi.org/10.1002/mar.21706>
- Zhang, et al. (2024). A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*. <https://doi.org/10.1145/3703155>