

Social Media Marketing and Online Purchase Intentions among Young Consumers

Mona Gupta

Guest Lecturer

Department of Commerce

Government Shaheed Kaushal Yadav College

Gunderdehi, Balod, Chhattisgarh



<http://www.gjirp.org/> || Vol. 2 No. 3 (2026): Special August Issue

Date of Submission: 15-07-2026

Date of Acceptance: 31-07-2026

Date of Publication: 20-08-2026

Abstract— Social media has evolved from an interpersonal communication environment into an algorithmically mediated marketplace in which young consumers encounter personalized advertisements, influencer content, electronic word-of-mouth, and interactive brand communication. This study investigates how social media marketing characteristics shape online purchase intentions among young consumers while distinguishing persuasive engagement from algorithmic personalization and perceived intrusiveness. A central research gap is identified in the limited integration of conventional social media marketing attributes with AI-enabled personalization, advertising transparency, privacy concern, and cross-platform behavioral signals within a single predictive framework. The proposed research conceptualizes purchase intention as the outcome of an interaction among content informativeness, entertainment value, social interactivity, electronic word-of-mouth, perceived personalization, brand trust, and privacy-related resistance. A hybrid analytical framework combining behavioral survey variables with machine-learning classification is proposed to identify both explanatory relationships and nonlinear patterns of purchase propensity. **Keywords—** Social Media Marketing, Online Purchase Intention, Young Consumers, AI Personalization, Brand Trust, Privacy Concern

Introduction

Social media has substantially altered the architecture of digital consumer markets. Platforms initially designed for interpersonal communication now combine social

networking, entertainment, advertising, influencer communication, product discovery, recommendation systems, and increasingly direct commerce functionality. Consequently, consumers do not necessarily enter a social platform with an explicit intention to purchase. Commercial stimuli are embedded within ordinary content consumption, and purchase consideration can emerge during browsing, interaction with creators, exposure to electronic word-of-mouth (eWOM), or repeated encounters with algorithmically selected brand messages. This distinction is particularly important when studying young consumers, whose product discovery processes increasingly occur within interactive and highly personalized digital environments.

Traditional digital advertising generally involved identifiable promotional placements presented to relatively broad audience segments. Contemporary social media marketing operates differently. Recommendation algorithms, behavioral profiling, predictive analytics, automated content selection, influencer ecosystems, and AI-supported personalization allow promotional messages to be matched with inferred consumer interests. As a result, two individuals using the same platform may encounter markedly different commercial environments. The marketing stimulus is therefore no longer determined exclusively by what a brand publishes; it is also influenced by the algorithmic processes governing what individual users actually see.

Recent scholarship confirms that social media marketing activities can influence purchase intentions through psychological and relational mechanisms rather than through a simple direct advertising effect. Koay, Cheah, and Goon (2023), for example, examined perceived social media

marketing activities through the Stimulus–Organism–Response perspective and demonstrated the importance of mechanisms involving brand trust and brand experience in explaining purchase intention. Such evidence suggests that the effectiveness of social media communication depends substantially on how consumers interpret and internalize marketing stimuli.

Consumer–brand interaction is another important component. Research published in the *Journal of Retailing and Consumer Services* indicates that effective social-media-based consumer–brand interactions can enhance purchase intention, while privacy concerns remain an important vulnerability within this relationship. These findings become increasingly relevant as platforms and advertisers rely on behavioral information to determine content relevance.

The development of artificial intelligence introduces another layer of complexity. AI-supported marketing systems can process behavioral signals, previous engagement, browsing activity, contextual information, and inferred preferences to improve advertisement targeting and content recommendations. Evidence published in *Computers in Human Behavior* indicates that AI-driven personalization can enhance perceived usefulness and trust in social media engagement, although privacy concerns can constrain its effectiveness. Similarly, recent research on personalized digital advertising has connected perceived personalization with purchase intention through relevance, trust, and usefulness.

The relationship between personalization and consumer response, however, should not be treated as uniformly positive. A highly relevant advertisement may reduce search costs and help consumers discover products consistent with their preferences. The same advertisement can produce discomfort when consumers perceive that the targeting process depends on excessive surveillance or unexpectedly detailed knowledge about their behavior. Research published in *Psychology & Marketing* in 2026 specifically examines consumers' awareness of digital tracking in personalized social media advertising and its relationship with trust and purchase-related outcomes. This creates an important conceptual distinction between **perceived relevance** and **perceived intrusiveness**.

Young consumers provide an especially important population for examining this tension. They frequently use social media for entertainment, social communication, product information, peer recommendations, creator content, and brand interaction. Consequently, purchase intention may emerge through multiple overlapping mechanisms. A product can first be encountered through a creator, reinforced through peer-generated content, presented again through personalized advertising, and subsequently evaluated using comments or

reviews. Conventional linear models of advertising exposure may therefore inadequately represent contemporary digital purchase journeys.

Research Gap

Existing research has made substantial progress in establishing relationships among social media marketing activities, customer engagement, brand trust, electronic word-of-mouth, and purchase intention. For example, previous studies have examined sequential mediation between social media marketing activities and purchase intention, customer engagement as an intervening mechanism, and consumer–brand interactions alongside privacy concerns. More recent work has separately investigated AI-driven personalization and consumer responses.



Fig. 1: Social media marketing approach for eCommerce companies

Source: <https://www.projectguru.in/influence-of-social-media-on-the-buying-behaviour-of-millennials/>

A narrower gap remains at the intersection of these research streams. Much of the literature relies on conventional structural models in which social media marketing or personalization is represented as an aggregate predictor. Less attention has been given to whether **different combinations of content quality, eWOM, interaction, AI personalization, trust, and privacy concern generate distinguishable high- and low-purchase-intention profiles among young consumers.**

Furthermore, personalization is commonly operationalized as a beneficial marketing characteristic. This approach may

overlook a nonlinear threshold: moderate personalization could improve relevance, whereas very intensive personalization combined with high tracking awareness could increase perceived intrusiveness and weaken purchase intention. The proposed study therefore treats algorithmic personalization as a potentially dual-effect mechanism rather than an automatically positive stimulus.

Accordingly, the manuscript is structured around the following research objective:

To develop and evaluate an interpretable machine-learning-assisted behavioral framework that predicts online purchase intention among young consumers from social media marketing attributes while determining how brand trust and privacy-related resistance alter the effectiveness of AI-driven personalization.

This perspective creates a bridge between consumer-behavior research and AI-based predictive analytics. Instead of asking only whether social media marketing statistically influences purchase intention, the proposed study asks **which combination of marketing and psychological signals best distinguishes consumers with different levels of purchasing propensity.**

Literature Review

Social Media Marketing as a Multidimensional Consumer Stimulus

Social media marketing is broader than the placement of advertisements on networking platforms. It includes brand-generated content, entertainment, information provision, customization, interactive communication, trend-oriented content, influencer collaboration, and mechanisms facilitating consumer-to-consumer information exchange. These characteristics transform marketing from predominantly one-directional communication into a dynamic process in which users can respond, redistribute, reinterpret, and publicly evaluate brand messages.

From the Stimulus–Organism–Response (S–O–R) perspective, marketing characteristics constitute environmental stimuli, psychological states represent the organism, and behavioral responses include engagement, product evaluation, and purchase intention. Contemporary empirical research has applied this logic to social media marketing. Koay et al. found that the influence of perceived social media marketing activities on purchasing outcomes can operate through relational mechanisms including brand trust and brand experience rather than functioning only as a direct pathway.

For young consumers, four marketing dimensions are particularly relevant: informativeness, entertainment,

interactivity, and personalization. Informativeness reduces uncertainty by providing product characteristics, prices, benefits, comparisons, or usage information. Entertainment increases hedonic value and can extend attention to branded material. Interactivity creates opportunities for consumers to communicate with brands and other users. Personalization increases the perceived correspondence between content and individual interests.

The proposed study extends this multidimensional interpretation by separating **human-perceived social marketing quality** from **algorithmically perceived personalization**. This separation is necessary because consumers may appreciate the content itself while simultaneously objecting to how the platform determined that the content should be shown to them.

Electronic Word-of-Mouth and Social Validation

Electronic word-of-mouth has become integral to online product evaluation. Comments, ratings, reviews, reposts, user-generated demonstrations, creator recommendations, and peer discussions provide information beyond formal brand communication. In contrast with traditional advertising, eWOM can function as a decentralized credibility mechanism because information originates from other consumers or apparently independent content creators.

Social validation can be especially influential when consumers face uncertainty regarding unfamiliar brands. High levels of positive interaction may operate as heuristic evidence that a product is accepted by other members of the digital community. Nevertheless, the effectiveness of eWOM depends on perceived credibility. Artificial engagement, undisclosed sponsorship, manipulated reviews, and commercially motivated influencer communication can weaken informational value.

This suggests that eWOM volume should not automatically be equated with persuasive strength. The present research therefore conceptualizes **eWOM credibility** as more relevant than mere exposure frequency. A consumer repeatedly exposed to promotional comments may exhibit lower purchase propensity than another consumer who encounters fewer but more credible peer recommendations.

Engagement, Brand Interaction, and Purchase Intention

Social media engagement describes cognitive, emotional, and behavioral involvement with content, brands, or communities. Observable engagement can include liking, sharing, saving, commenting, following, or clicking. Psychological engagement, however, may exist without visible interaction. A consumer can carefully examine a product video without liking or commenting on it.

Evidence indicates that social interaction can support purchase intention. Research examining consumer–brand interactions has shown that stronger social media interaction can contribute to purchasing outcomes, although privacy remains a critical concern. Other empirical work has also examined customer engagement as an intermediary between social media marketing activities and purchase intention.




Fig. 2: Research Model of the Impact of Social Media on Consumer Purchase Intentions

Source: https://www.researchgate.net/figure/Research-Model-of-the-Impact-of-Social-Media-on-Consumer-Purchase-Intentions_fig1_343296536

These findings support the inclusion of engagement in purchase-intention models but also highlight a measurement challenge. Visible engagement metrics alone may overestimate actual buying propensity. Users sometimes interact with humorous or entertaining branded content without considering the advertised product. The proposed framework therefore distinguishes **content engagement** from **commercial intention**, preventing engagement from being interpreted as a direct substitute for purchasing propensity.

AI-Driven Personalization and Perceived Relevance

AI-based personalization represents one of the most significant changes in contemporary social media marketing. Machine-learning systems can infer consumer interests from previous interactions and dynamically rank advertisements, creators, products, and promotional content. Personalization can therefore reduce information overload by prioritizing material predicted to be relevant to a specific user.

Recent evidence demonstrates the importance of this mechanism. Teepapal's research in *Computers in Human Behavior* examined AI-driven personalization in social media engagement and highlighted perceived usefulness and trust while also identifying privacy concerns as a challenge for

personalized marketing. Research on AI-powered personalized advertising has likewise examined the roles of relevance, usefulness, and trust in explaining purchase intention.

The implication is that personalization is unlikely to influence purchasing simply because an algorithm selected an advertisement. The advertisement must first be perceived as relevant and useful. Thus, the proposed framework conceptualizes **perceived relevance as the cognitive benefit of personalization**, whereas perceived intrusiveness represents its potential psychological cost.

Brand Trust as a Conversion Mechanism

Purchase intention requires consumers to accept a degree of uncertainty. Online consumers may be unable to physically inspect a product, directly evaluate the seller, or verify every promotional claim. Brand trust therefore reduces perceived transactional uncertainty.

Empirical research supports this role. Koay et al. demonstrated that brand trust forms part of the mechanism connecting social media marketing activities with subsequent brand experience and purchase intention. Research examining brand trust as a mediator has similarly reported its relevance to relationships between social media marketing activities and purchase intention.

In an AI-mediated environment, trust has at least two distinguishable targets: **trust in the brand** and **trust in the personalization process**. A consumer may trust a retailer's products but distrust the platform's data collection practices. Conversely, a user may regard recommendations as useful while remaining uncertain about an unfamiliar seller. Combining these dimensions into one generalized trust variable can conceal important behavioral differences.

The proposed study consequently emphasizes brand trust while separately measuring privacy and personalization perceptions. This design should provide greater interpretability regarding whether purchasing resistance arises from doubts about the seller or discomfort with the targeting mechanism.

Privacy Concern, Tracking Awareness, and the Personalization Paradox

Personalization requires information. The more precisely a marketing system adapts content to individual preferences, the more likely consumers are to infer that behavioral information has been collected or analyzed. This creates what may be described as a personalization paradox: consumers value relevance while simultaneously expressing concern about surveillance and data use.

Recent evidence makes this issue particularly salient. The 2026 *Psychology & Marketing* study by Tangari investigates how awareness of tracking used for personalized social media advertising shapes trust, purchase intention, and willingness to pay. Research on retail social media interactions similarly identifies privacy as a persistent vulnerability even when consumer–brand interaction is otherwise effective.

The present manuscript therefore proposes that privacy concern should not be treated merely as a demographic control variable. Instead, it represents a **behavioral resistance signal** capable of changing how personalization translates into purchasing propensity.

A consumer with low privacy concern may interpret a highly personalized advertisement as convenient. Another consumer exposed to the identical level of personalization may interpret it as intrusive. Consequently, the marginal effect of personalization may vary considerably across consumers.

Machine Learning and the Need for Nonlinear Purchase-Intention Modelling

Much of the consumer-behavior literature relies on regression or structural equation modeling. These approaches remain valuable for hypothesis testing and explaining directional relationships. However, the contemporary social media environment may contain nonlinearities and interactions that are difficult to represent through a single linear specification.

For example, entertainment value may increase purchase intention only when content is sufficiently informative. High personalization may be effective when trust is high but counterproductive when privacy concern is elevated. Likewise, strong eWOM may compensate for low familiarity with a brand but contribute little when consumers already possess high brand trust.

The proposed study therefore introduces an interpretable machine-learning layer rather than replacing theory-driven statistical analysis. Gradient-boosted decision trees will be compared with regularized logistic regression and a compact multilayer perceptron for classifying purchase-intention propensity. Explainability techniques such as SHAP-based feature attribution will subsequently be used to determine whether personalization, trust, eWOM credibility, interactivity, informativeness, or privacy concern contributes most strongly to individual predictions.

This methodological orientation differentiates the research from conventional studies that report only path coefficients. The intended contribution is not merely higher predictive accuracy but the identification of **behavioral configurations** associated with purchase propensity.

Conceptual Synthesis

The literature indicates that social media marketing affects consumer decision-making through interconnected informational, emotional, social, and technological mechanisms. Informativeness and eWOM reduce informational uncertainty; entertainment and interactivity support attention and engagement; personalization increases relevance; brand trust reduces transactional uncertainty; and privacy concern can introduce resistance to algorithmically targeted persuasion.

The central proposition emerging from this review is therefore that online purchase intention among young consumers cannot be adequately represented as a simple function of social media exposure. It is more appropriately conceptualized as a conditional behavioral outcome:

Proposed Methodology

The study adopts a **quantitative, cross-sectional, predictive-explanatory research design** to examine how social media marketing characteristics influence online purchase intentions among young consumers. Unlike conventional studies that depend exclusively on structural equation modelling, the proposed framework combines psychometric measurement with machine-learning classification. This design permits two complementary analyses: identification of theoretically meaningful determinants of purchase intention and prediction of consumers displaying relatively high purchasing propensity.

The target population comprises social media users aged **18–30 years** who have encountered product-related content on platforms such as Instagram, YouTube, Facebook, Snapchat, and short-form video or social-commerce environments during the preceding six months. A proposed sample of **480 respondents** is used for the empirical framework. Respondents are recruited using stratified purposive sampling to obtain variation in age, gender, educational background, frequency of social media use, and online shopping experience.

The conceptual framework contains eight principal predictor constructs: social media informativeness (INF), entertainment value (ENT), interactivity (INT), eWOM credibility (EWC), perceived personalization (PER), perceived relevance (REL), brand trust (BT), and privacy concern/perceived intrusiveness (PRI). Online purchase intention (OPI) constitutes the outcome construct.

Responses are collected through a structured questionnaire using five-point Likert-type items ranging from 1 = strongly disagree to 5 = strongly agree. Screening questions ensure that participants are active social media users and have previously encountered commercial content through social platforms.

The methodological innovation is the explicit separation of **personalization benefit** from **personalization resistance**. Perceived relevance measures whether algorithmically selected commercial content corresponds with users' interests, whereas privacy concern captures discomfort regarding behavioral tracking and personal-data-driven targeting.

For predictive modelling, respondents are divided into two purchase-intention categories. Participants whose composite OPI scores exceed a predetermined scale threshold are classified as **high purchase-intention consumers (Class 1)**, while the remaining participants form the **low/moderate purchase-intention group (Class 0)**. As a robustness check, the threshold should also be varied in sensitivity analysis rather than selected after observing model performance.

The following conceptual relationships guide the analysis:

H1: Social media informativeness positively contributes to online purchase intention.

H2: Entertainment value positively contributes to online purchase intention.

H3: Interactive social media marketing positively contributes to online purchase intention.

H4: Credible electronic word-of-mouth positively contributes to online purchase intention.

H5: Perceived personalization positively contributes to purchase intention primarily when it generates perceived relevance.

H6: Brand trust positively contributes to online purchase intention.

H7: Privacy concern and perceived intrusiveness negatively contribute to online purchase intention.

H8: Privacy concern weakens the positive contribution of AI-driven personalization to online purchase intention.

Rather than assuming that every relationship is linear, the machine-learning component evaluates interactions and threshold effects among these variables.

Experimental Setup

Dataset

The proposed experimental dataset contains **480 valid consumer observations**. Each observation represents one respondent rather than an individual social media post. Approximately 34 questionnaire indicators are used to construct the principal variables.

The dataset contains four categories of information:

1. **Consumer profile:** age, gender, education and online-shopping experience.
2. **Platform behavior:** daily social media exposure, product-discovery frequency, engagement with commercial content and previous social-media-influenced purchases.
3. **Marketing perceptions:** informativeness, entertainment, interactivity, eWOM credibility, personalization and relevance.
4. **Psychological responses:** brand trust, privacy concern, perceived intrusiveness and purchase intention.

The survey should not collect passwords, private messages, exact browsing histories, or personally identifying platform credentials. This privacy-preserving approach is particularly appropriate because the study itself investigates consumer perceptions of behavioral targeting.

Data Preprocessing

Data preprocessing proceeds through several stages. First, incomplete responses with excessive missing information are removed. Limited item-level missingness is handled using median imputation within the relevant construct rather than replacing observations using the global dataset mean.

Second, reverse-coded questionnaire items are transformed before calculating composite construct scores. Internal consistency is evaluated before composite variables are created. Construct scores are subsequently standardized for models sensitive to variable scale.

Third, categorical demographic variables are transformed using one-hot encoding. Continuous variables are standardized using parameters estimated exclusively from the training data to prevent information leakage.

Fourth, the dataset is stratified into **70% training, 15% validation, and 15% independent test subsets**. Class distributions are maintained as consistently as possible across the partitions.

Potential class imbalance is examined after defining the purchase-intention outcome. If substantial imbalance is present, class-weighted learning is preferred as the primary strategy. Synthetic oversampling, if tested, is restricted to training folds and never applied before dataset partitioning.

Model Architecture

Three predictive models are proposed.

Regularized Logistic Regression (LR) provides an interpretable linear baseline. L2 regularization controls coefficient instability arising from correlated marketing variables.

Gradient-Boosted Decision Trees (GBDT) constitute the principal model. Sequentially constructed shallow decision trees allow nonlinear relationships and interactions to be captured without requiring the researcher to specify every interaction manually. This architecture is particularly appropriate for tabular consumer survey data.

Multilayer Perceptron (MLP) provides a compact neural-network benchmark. The network contains an input layer corresponding to the processed predictor variables, two hidden layers of 64 and 32 neurons, ReLU activation, dropout regularization, and a sigmoid output neuron for binary classification.

The objective is not to assume that a neural architecture will outperform other algorithms. Instead, the comparison tests whether the additional flexibility of neural modelling provides meaningful predictive improvement over a tree ensemble and an interpretable linear baseline.

Training Strategy

Model development uses only the training and validation partitions. Logistic-regression regularization strength, GBDT learning rate, tree depth, number of estimators and MLP learning parameters are selected through validation-based tuning.

For the MLP, binary cross-entropy is employed as the loss function with Adam optimization. Early stopping is used when validation loss fails to improve over consecutive epochs, reducing unnecessary training and limiting overfitting.

The final selected configurations are then evaluated once against the untouched test set.

To strengthen robustness, five-fold stratified cross-validation is additionally conducted on the development data. Model selection prioritizes **ROC-AUC** and **F1-score** rather than accuracy alone. ROC-AUC summarizes discrimination across classification thresholds, while F1 combines precision and recall. These are established classification-evaluation measures.

Evaluation Metrics

Five principal metrics are used:

Accuracy measures the overall proportion of correctly classified consumers.

Precision indicates the proportion of predicted high-intention consumers who actually belong to the high-intention category.

Recall evaluates the proportion of actual high-intention consumers successfully identified by the model.

F1-score provides the harmonic balance between precision and recall.

ROC-AUC measures the ability of the model to rank positive observations above negative observations across decision thresholds.

Because predictive performance alone cannot explain consumer behavior, the best-performing model is interpreted using **SHAP (SHapley Additive exPlanations)**. Lundberg and Lee introduced SHAP as a unified additive feature-attribution framework assigning importance values to features for individual predictions.


Results and Discussion

Important empirical-status note: because no real respondent dataset was supplied with the manuscript, the numerical results below are **simulated/proposed experimental results for manuscript design**. They must not be represented as observations from an actually conducted survey. For journal submission, they should be replaced by outputs calculated from the collected dataset.

The illustrative experiment indicates that the gradient-boosted model provides the strongest overall discrimination between high- and lower-purchase-intention consumers. Its advantage over logistic regression suggests that social-media-mediated purchasing behavior may involve nonlinear relationships that cannot be completely represented through additive linear coefficients.

Table 1. Illustrative Comparative Performance of Purchase-Intention Prediction Models

Predictive Model	Accuracy	Precision	Recall	F1-Score	ROC - AUC
Regularized Logistic Regression	0.764	0.748	0.781	0.764	0.817
Multilayer Perceptron	0.806	0.798	0.812	0.805	0.862
Gradient-Boosted	0.833	0.821	0.846	0.833	0.891



The illustrative ROC-AUC of 0.891 for GBDT indicates stronger ranking performance than the other candidate models. The difference between GBDT and logistic regression is theoretically meaningful because the proposed framework assumes that consumer response to personalization is conditional rather than uniformly positive.

SHAP-based interpretation of the simulated model identifies **brand trust, perceived relevance, eWOM credibility, perceived personalization, and privacy concern** as the most influential variables. However, their effects differ substantially.

Brand trust emerges as the strongest positive predictor. This result is consistent with the argument developed in Part 1: social media exposure does not automatically translate into purchasing propensity. Consumers must first consider the seller or brand sufficiently credible to reduce transactional uncertainty.

Perceived relevance provides another strong positive contribution. Consumers who believe that promoted products correspond with their interests are more likely to enter the high-purchase-intention category. This supports the commercial value of personalization but also clarifies that the relevant psychological mechanism is not personalization itself. **Personalization creates value when consumers experience it as relevance.**

Electronic word-of-mouth credibility also makes a substantial contribution. Positive reviews and recommendations appear most useful when consumers consider them authentic. The result reinforces the distinction between engagement volume

and credible social endorsement. Thousands of visible reactions cannot necessarily substitute for trusted consumer information.

A particularly important pattern concerns privacy concern. SHAP interaction analysis in the illustrative experiment indicates that personalization contributes positively when privacy concern is low or moderate. At higher privacy-concern levels, however, the marginal contribution of personalization declines and can become negative.

This pattern supports the proposed **personalization–intrusiveness threshold**. A young consumer may initially appreciate recommendations that correspond with prior interests. When targeting becomes unexpectedly specific, repeated, or apparently dependent on extensive behavioral tracking, the same personalization mechanism may trigger resistance.

Entertainment value displays a smaller direct predictive contribution than brand trust and relevance. This does not imply that entertainment is unimportant. Rather, entertaining content may attract attention and stimulate engagement without necessarily generating commercial intent. This distinction is valuable for marketers because engagement statistics such as likes, shares, and views can exaggerate actual purchasing potential.

Interactivity exhibits a moderate positive contribution. Responsive brand communication, accessible product information, and opportunities for consumers to ask questions appear more closely connected to purchasing propensity than superficial interactions alone.

The findings therefore support a layered interpretation of social media marketing:

Attention generation → perceived relevance → credibility assessment → brand trust → purchase propensity

with privacy-related resistance capable of interrupting the pathway.

From a marketing perspective, this means that maximizing advertisement exposure or personalization intensity is not necessarily optimal. AI-assisted social media strategies should optimize for **relevant, credible and non-intrusive personalization** rather than maximum behavioral targeting.

The model comparison also has methodological implications. Logistic regression remains useful because coefficient direction and magnitude are straightforward to communicate. The MLP demonstrates the potential of neural representation but offers only moderate improvement in the illustrative experiment. GBDT achieves the strongest balance between predictive capacity and post-hoc interpretability.

SHAP analysis is particularly valuable in this context because it allows researchers to move from the question “Which model predicts purchase intention most accurately?” to the more substantive question “Which consumer and marketing characteristics pushed an individual prediction toward high or low purchase intention?” The original SHAP framework was specifically developed to provide feature-level explanations for complex model predictions.

Future Scope

Future studies should extend the framework from purchase **intention** to verified behavioral conversion. With appropriate informed consent and privacy safeguards, researchers could compare survey responses with anonymized click-through, shopping-cart, or purchase-event information.

A longitudinal design would provide a stronger understanding of temporal effects. Consumers could be observed at several intervals to determine whether repeated personalization gradually increases familiarity or instead produces advertising fatigue and privacy resistance.

Another direction involves multimodal AI. Future models could analyze characteristics of advertising images, video structure, caption sentiment, linguistic style, creator credibility, and consumer responses simultaneously. Such analysis could determine whether particular combinations of visual and textual cues generate stronger purchasing responses.

Cross-cultural research is also necessary. Privacy expectations, influencer credibility, digital-payment adoption, and trust in online sellers vary between markets. Comparative studies could determine whether the personalization–intrusiveness threshold remains stable across cultural and regulatory environments.

Future experiments could additionally introduce causal designs in which participants are randomly exposed to low, moderate, and high personalization conditions. Such experiments would provide stronger evidence regarding whether excessive personalization actually causes resistance.

Finally, explainable AI should remain central to consumer-prediction research. High predictive accuracy without an understanding of the underlying drivers provides limited theoretical contribution. Feature-attribution techniques can help connect machine-learning predictions with established consumer-behavior theory.

Conclusion

This study develops a distinct framework for understanding **social media marketing and online purchase intentions among young consumers** by integrating traditional

marketing constructs with AI-enabled personalization, privacy concern, and interpretable machine learning.

The central argument is that contemporary social media marketing cannot be adequately understood through exposure or engagement alone. Young consumers operate within algorithmically curated commercial environments in which informativeness, entertainment, interactivity, eWOM, personalization, relevance, trust, and privacy concerns interact in determining purchasing propensity.

The proposed comparative experiment evaluates regularized logistic regression, a multilayer perceptron, and gradient-boosted decision trees. Illustrative results identify GBDT as the strongest predictive architecture, while SHAP analysis provides a mechanism for explaining individual predictions. SHAP is grounded in the feature-attribution framework introduced by Lundberg and Lee.

The study's most important conceptual contribution is the **personalization–intrusiveness threshold**. AI-driven personalization should not be assumed to exert an indefinitely positive effect. Personalization appears commercially useful when it produces relevance, but its benefit may diminish when consumers infer excessive tracking or experience targeted advertising as intrusive.

Accordingly, effective social media marketing should prioritize trusted, informative, relevant, and privacy-sensitive consumer experiences rather than simply increasing targeting intensity. This interpretation positions online purchase intention as the product of an interaction between marketing quality, algorithmic relevance, social credibility, consumer trust, and resistance to perceived surveillance.

The framework contributes to both marketing analytics and AI-based consumer-behavior research by demonstrating how interpretable machine learning can supplement theory-driven analysis without replacing it. More importantly, it provides a research design in which predictive performance and behavioral explanation are evaluated simultaneously, creating a stronger foundation for future empirical research on AI-mediated social commerce.

References

- [1] K. Y. Koay, C. W. Cheah, and S. W. U. Goon, “How Do Perceived Social Media Marketing Activities Foster Purchase Intentions? A Multiple Sequential Mediation Model,” *Journal of Global Marketing*, vol. 36, no. 3, pp. 210–224, 2023, doi: 10.1080/08911762.2023.2207072.
- [2] K. Y. Koay, C. W. Cheah, and G. L. Goon, “Do brands’ social media marketing activities matter? A moderation analysis,” *Journal of Retailing and Consumer Services*, vol. 64, p. 102794, 2022, doi: 10.1016/j.jretconser.2021.102794.

[3] T. Teepapal, "AI-driven personalization: Unraveling consumer perceptions in social media engagement," *Computers in Human Behavior*, vol. 165, p. 108549, 2025, doi: 10.1016/j.chb.2024.108549.

[4] M. S. Alzaidi and G. Agag, "The role of trust and privacy concerns in using social media for e-retail services: The moderating role of COVID-19," *Journal of Retailing and Consumer Services*, vol. 68, p. 103042, 2022, doi: 10.1016/j.jretconser.2022.103042.

[5] E. Aguirre, D. Mahr, D. Grewal, K. de Ruyter, and M. Wetzels, "Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness," *Journal of Retailing*, vol. 91, no. 1, pp. 34–49, 2015, doi: 10.1016/j.jretai.2014.09.005.

[6] T. A. Ngo, C. T. Bui, H. K. L. Chau, and N. P. N. Tran, "Electronic word-of-mouth (eWOM) on social networking sites (SNS): Roles of information credibility in shaping online purchase intention," *Heliyon*, vol. 10, no. 11, p. e32168, 2024, doi: 10.1016/j.heliyon.2024.e32168.

[7] N. Ameen, A. Tarhini, A. Reppel, and A. Anand, "Customer experiences in the age of artificial intelligence," *Computers in Human Behavior*, vol. 114, p. 106548, 2021, doi: 10.1016/j.chb.2020.106548.

[8] M. Ashfaq, J. Yun, S. Yu, and S. M. C. Loureiro, "I, Chatbot: Modeling the determinants of users' satisfaction and continuance intention of AI-powered service agents," *Telematics and Informatics*, vol. 54, p. 101473, 2020, doi: 10.1016/j.tele.2020.101473.

