Drivers of Generative AI Acceptance in Agriculture: An AHP Study with Synthetic Experts*

Murat Tahir Çaldağ

Abstract— **Although transforming agriculture with generative artificial intelligence (AI) can provide significant improvements, the acceptance of these tools is underexplored. Therefore, in this study the drivers of generative AI acceptance in agriculture are explored using a systematic literature review (SLR) and an analytic hierarchy process (AHP) analysis. SLR results presented 5 main drivers and 15 sub-drivers that are mapped to the proposed acceptance framework. The AHP analysis of the proposed framework is conducted with synthetic experts generated with generative AI to provide an early exploratory analysis of the research. The findings presented perceived ease of use as the main driver followed by social influence while perceived risk and benefits ranked lowest in generative AI acceptance in agriculture. Accessibility, financial resources and culture were ranked highest sub-drivers, while traceability, hallucination and simplicity were the lowest. The results emphasized ease of use as the priority factor in farmers' acceptance decisions regarding generative AI in agriculture. The contribution of this study is to present an extended generative AI acceptance model based on TAM and UTAUT frameworks that are preliminarily assessed with synthetic expert based AHP analysis.**

I. INTRODUCTION

The rapid advancement of generative artificial intelligence is transforming traditional industries through new ways of automation, interaction and decision-making. While industries such as healthcare, education and software engineering highly integrate generative AI, agriculture is in its early stages of use and acceptance of generative AI. Even though the potential benefits such as assisting farmers with information retrieval, diagnosing crop issues and generating care advice, the acceptance of generative AI in agriculture remains surprisingly underexplored.

Existing literature on generative AI practices in agriculture mainly focus machine learning for precision agriculture practices or accuracy evaluation of information retrieval while the acceptance of generative AI tools are scarce. Farmers' perceptions, trust, and social contexts are critical drivers of technology acceptance. Therefore, understanding how farmers evaluate the usefulness, ease of use, social influence, facilitating conditions, and perceived risks of generative AI is essential for guiding policy, design, and the implementation of AI-driven agricultural solutions.

*There is no funding for this Research.

Murat Tahir Çaldağ is with the Department of Technology and Knowledge Management, Başkent University, Ankara, Turkey (e-mail: mtcaldag@baskent.edu.tr)

To address this research gap, this study investigates generative AI acceptance in agriculture. First a systematic literature review is conducted for identifying the drivers of generative AI acceptance in agriculture. The identified drivers are mapped to an acceptance framework extended by Technology Acceptance Model (TAM) and Unified Theory of Acceptance and Use of Technology (UTAUT) for agriculture with the addition of perceived risk dimension. To assess the framework an Analytic Hierarchy Process (AHP) analysis using judgments from synthetic experts is conducted. The contribution of this hybrid approach presents an understanding of generative AI acceptance in agriculture and demonstrates the feasibility of synthetic expert modelling and usage for the early-stage exploration of decision frameworks.

The second section provides details on SLR of generative AI acceptance in agriculture. The third section presents the proposed framework. The fourth section explains synthetic expert generation and the AHP analysis. The fifth section presents the research findings. The last section emphasizes the study's contributions, limitations, and results.

II. SYSTEMATIC LITERATURE REVIEW

The generative AI adoption throughout the agriculture sector is assessed by a SLR based on Kitchenham (2004). The search was carried out in Scopus and Web of Science databases with the established search string: ("generative artificial intelligence" OR "generative ai" OR "GenAI" OR "large language model" OR "LLM" OR "chatbot" OR "chatGPT") AND ("agriculture" OR "farm" OR "crop" OR "cultivation" OR "livestock" OR "ranch" OR "horticulture" OR "aquaculture" OR "irrigation" OR "agronomy")). The inclusion and exclusion criteria of the search consisted of date, publication type, and language criteria. The date criterion included publications published after 2015 to exclude irrelevant studies. The language of the search was carried out in English to generate a holistic review of the literature. Only articles and proceeding papers were included in the search to provide significant and peer-reviewed studies.

The initial search presented 1130 academic studies that were filtered based on the exclusion criteria to 911 articles and proceedings on generative AI in the context of agriculture. After removing duplicate studies and evaluating relevance based on title, abstract, and keywords, 200 studies remained. Lastly the relevancy is checked by full-paper analysis which presented 11 related papers on generative AI acceptance in agriculture that is given in Table 1. Even though the primary pool of SLR included several studies on the context of generative AI and agriculture, there is a research gap in the acceptance frameworks, models and evaluations. Therefore, an acceptance framework is proposed to evaluate farmers' attitudes and the use of generative AI.

Table 1: Primary Pool of SLR.

Reference	*Aim of the Study*
[2]	To test an LLM-based virtual assistant in Nigeria for farmers to retrieve information on indigenous vegetables.
[3]	To evaluate an agent-based LLM on data sharing across agriculture IoT platforms.
[4]	To design a multilingual, voice-enabled chatbot to assist Indian farmers.
[5]	To extract information on crop breeding research and programs.
[6]	To develop a system to classify bonsai species and generate care advice.
[7]	To integrate a crop recommendation engine and a question-answering chatbot to guide farmers.
[8]	To explore how GenAI tools can expand livelihood opportunities with community members in a rural Indian village
[9]	To propose a chatbot framework to support farmers in India.
[10]	To develop a documentation-driven chatbot to assist Peruvian dairy producers with feeding and animal-care decisions.
[11]	To evaluate GPT4 for improved decision support on broiler producers in Brazil.
[12]	To examine the predictors of AI-chatbot adoption among rural women farmers in Zimbabwe.

III. PROPOSED FRAMEWORK

To address the research gap on generative AI acceptance in agriculture, existing models on acceptance and adoption can provide significant guidance. The commonly used frameworks on acceptance and adoption of technologies are Technology Acceptance Model (TAM), Unified Theory of Acceptance and Use of Technology (UTAUT), Theory of Planned Behavior (TBP), Technology Organization Environment (TOE) and Diffusion of Innovation (DOI) [13], [14], [15], [16], [17]. While TAM, UTAUT and TBP models' assessment of individual acceptance TOE and DOI models evaluate the organizational adoption of technologies. In this study, the individual acceptance of generative AI by farmers is analyzed by extending the TAM and UTAUT frameworks.

The proposed framework is an extension of TAM's main drivers, perceived usefulness and perceived ease of use, with UTAUT's facilitating conditions and social influence. Another significant extension is the perceived risk dimension, which emphasizes AI acceptance risk.

The acceptance drivers of generative AI in agriculture identified by the SLR results is categorized under five main categories given in Figure 1 as perceived usefulness, perceived ease of use, facilitating conditions, social influence and perceived risks. The subcategories of these drivers include accuracy, productivity, traceability, accessibility, scalability, simplicity, financial resources, knowledge & awareness, government support, culture, infrastructure, sustainability expectations, overreliance, bias, and hallucinations.

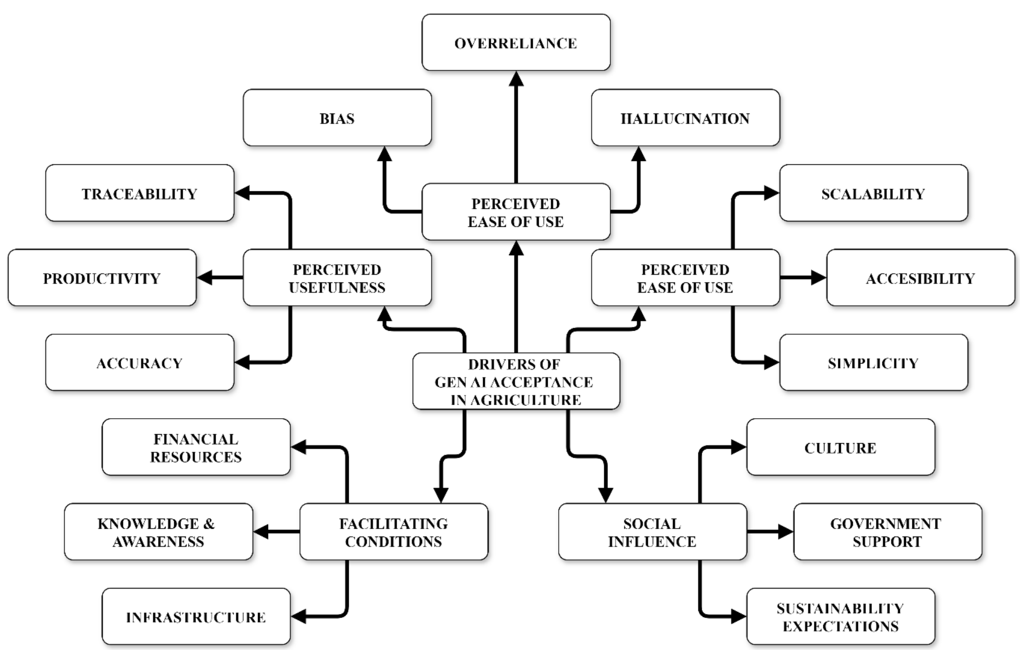

Figure 1: Proposed GenAI Acceptance Framework

Perceived usefulness in the context of the study is the degree farmers believe using generative AI would enhance their farming performance [14]. Accuracy is associated with the correctness and reliability of the output created by generative AI [5], [6]. Productivity is the perceived effectiveness of time or resources with the use of generative AI in farming tasks [4], [11]. Traceability refers to the transparency caused by the use of generative AI for improvements in record keeping and documentation [11].

Perceived ease of use refers to the degree of effort needed to use generative AI in agricultural tasks [14]. Scalability is the perceived effort in implementing generative AI across multiple agricultural tasks [6], [9]. Accessibility is associated with the degree of overcoming the barriers on using generative AI in farm tasks [8]. Simplicity is referred to the natural feeling while using generative AI in farming [12].

Facilitating conditions is defined as the farmers belief in the existing support and infrastructure to use generative AI [17]. Financial resources are defined as the farmers perception of the compensation of generative AI use and integration into agricultural tasks [2]. Knowledge and awareness refers to farmers perceived expertise on generative AI use [2]. Infrastructure is associated with availability of the internet, electricity and electronic devices required to integrate generative AI with agricultural tasks [2].

Social influence in the context of generative acceptance in agricultural practices is defined as the farmers perception to use generative AI in agricultural tasks based on peers and external pressures [17]. Culture refers to communities' resistance to generative AI use in agriculture [2]. Government support includes tax policies and regulatory practices on generative AI integration throughout farming communities [2]. Sustainability expectations refer to practices that promote self-sufficient and eco-friendly approaches to generative AI acceptance throughout farm communities [12].

Perceived risk is associated with the negative outcomes of using generative AI, which can hinder adoption in agricultural contexts [18]. Overreliance refers to excessive trust in generative AI in farming operations [18]. Bias is related to the generative AI's training data being discriminative or subjective [18]. Hallucination is the risk of creating false content with generative AI [19]. Sub-drivers derived from broader AI literature rather than agriculture-specific acceptance studies (overreliance, bias, hallucination) are considered theoretical extensions of the perceived risk dimension. The identified five main drivers and 15 sub-drivers of generative AI acceptance in agriculture require assessment to determine the consistency of the proposed framework. To assess the framework an AHP analysis is conducted with synthetic experts.

IV. METHODOLOGY

The research methodology consists of two stages. The first stage includes constructing a demographically grounded synthetic expert population using an LLM-based persona generation approach. The second stage consists of evaluating the proposed GenAI acceptance in agriculture framework with the judgements of synthetic experts generated. The use of synthetic experts in this study is not intended to replace empirical human judgment or to serve as a definitive validation of generative AI acceptance in agriculture. Instead, synthetic experts are employed as an exploratory and theory-building instrument to support early-stage analysis of complex socio-technical frameworks. Accordingly, the AHP results presented in this study should be interpreted as exploratory weight estimates rather than empirically validated priorities.

A. Synthetic Expert Generation

Synthetic experts are artificially created specialists that can provide insights just as human experts in a given domain of knowledge. The major benefit of using synthetic experts in research is providing preliminary insights into the subject before further analysis. The generation of synthetic experts with varied demographic attributes by large language models (LLM) is referred to as silicon sampling [20]. Silicon sampling has been shown to generate approximate population-level responses across various domains, though with important limitations [20], [21], [22].

However, synthetic expert generations can vary with minor changes in prompt wording, and the same prompt yields significantly different results across different time periods because of model updates [23]. To mitigate these risks a synthetic farmer generation prompt was designed according to four evidence-based principles such as using fixed demographic quotas, having internal consistency rules, adding an explicit attitude variable and having a structured output. Rather than randomizing the synthetic expert profiles, exact distributional quotas were specified for all key attributes (gender, region, farm size, digital literacy, and technology attitude) that supports the results to be more representative and reproducible outputs across model runs [24]. Explicit rules were specified linking demographic attributes to each other that prevents the model from generating conflicting personas [25]. Silicon sampling research has shown that LLMs tend toward socially desirable, neutral outputs when attitudinal variation is not explicitly specified [26]. Therefore, adding a tech attitude field with three levels and fixed distributional quotas forces diversity into the synthetic population. Finally, to address the variation in the output, a structured JSON output format was selected. Fixing a precise JSON output schema improves the reduction in variation across runs and enables repeatability [23].

The synthetic expert generation and expert judgements were conducted using ChatGPT (OpenAI, March 2026). As the platform dynamically assigns GPT-5 series models without exposing exact version identifiers to the user, a specific model version cannot be confirmed. The system employed a reasoning-capable model optimized for structured analytical tasks. To support reproducibility, the prompts used are provided as supplementary materials. To address the research gap on generative AI acceptance in agriculture 22 farmer profiles are generated. Table 2 presents 5 representative profiles selected to reflect diverse income, education and attitude levels.

Table 2: Representative Synthetic Farmer Profiles

ID	Education	Income	Farmer's Description
1	Primary	Low	He operates a small maize farm in rural Kenya with limited mechanization. He is cautious about digital technologies because unreliable internet and cost make them difficult to trust.
5	Master's or higher	Medium	He operates a large precision agriculture wheat farm in the United States using advanced machinery. He strongly supports AI tools because they optimize irrigation and yield forecasting.
6	Primary	Medium	She manages a small mixed farm in rural Jordan focused on vegetables and goats. She avoids new technologies because limited training and financial constraints make adoption risky.
14	Master's or higher	High	She manages a large technology-driven greenhouse farm in Japan. She supports AI systems because they enhance precision farming and resource efficiency.
22	High school	Low	She operates a poultry farm in Thailand supplying eggs to regional markets. She is interested in technology if it improves disease monitoring and farm productivity.

The full set of 22 profiles is available in the supplementary materials (https://doi.org/10.5281/zenodo.20044579). Synthetic farmers age, gender, education level, farm size, farm type, income level, digital literacy level, experience, region and a short description are generated with ChatGPT.

B. Analytic Hierarchy Process

To evaluate the proposed framework on generative AI acceptance in agriculture an AHP analysis is conducted. AHP analysis is generally used as a support for choosing alternatives and prioritization. AHP analysis is developed by [27] as a multi-criteria decision making method. The analysis includes creation of the hierarchical model, gathering experts' judgements on the model, and evaluating pairwise comparisons.

In this study the drivers of generative AI acceptance in agriculture are identified with SLR and are mapped with an extension on TAM and UTAUT models for further analysis. The AHP step of the study included generation and evaluation of the comparison matrices for the drivers' weights and significance. The AHP Judgements of the synthetic farmers were also generated by ChatGPT.

The comparison matrix (A) is created according to the number of factors (n) in each category given by (1). The inputs of the comparison matrix are obtained from a nine-point scale (Saaty, 1990).

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix} \qquad (1)$$

where $a_{ij} = \frac{1}{a_{ji}}$ for $i, j = 1, \ldots, n.$

The Consistency Ratio (CR) of the matrix A is calculated by following the steps given below.

Step 1. The weight of each factor (ω_i) is evaluated as

$$\omega_i = \sqrt[n]{a_{i1}a_{i2} \ldots a_{in}} \qquad i = 1, \ldots, n. \qquad (2)$$

Step 2. The normalized weight ($N\omega_i$) is computed from

$$N\omega_i = \frac{\omega_i}{\omega_1 + \omega_2 + \cdots + \omega_n} \qquad i = 1, \ldots, n. \qquad (3)$$

Step 3. The λ_{max} value is given by

$$A \quad \omega = \lambda_{max} \quad \omega \qquad (4)$$

Step 4. The CR is calculated by

$$CR = \frac{CI}{RI} \qquad (5)$$

where the consistency Index (CI) is

$$CI = \frac{\lambda_{max} - n}{n-1} \qquad (6)$$

and RI denotes the random index determined by [27] given in Table 3.

Table 3: AHP Random Index.

n	1	2	3	4	5
RI	0	0	0.58	0.9	1.21
n	**6**	**7**	**8**	**9**	**10**
RI	1.24	1.32	1.41	1.45	1.49

The CR value of the matrix requires to be less than 0.10 for reliability in factor comparison [27]. If the value of CR is greater than 0.10, the comparison matrix is unsuitable for analysis and will require new data from existing or new participants.

The practicality of the AHP analysis in this study is to provide an initial assessment of the research before investing more resources. The calculations and gathering of the comparison matrices are conducted by Microsoft Excel spreadsheets.

Table 4: AHP Analysis Results.

Main Driver	*Relative Weight*	*Sub-Driver*	*Local Weight*	*Local Rank*	*Global Weight*	*Global Rank*
Perceived Usefulness	0,1716	Accuracy	0,6513	1	0,1118	4
		Productivity	0,2404	2	0,0413	11
		Traceability	0,1083	3	0,0186	15
Perceived Ease of Use	0,2305	Accessibility	0,6180	1	0,1424	1
		Scalability	0,2364	2	0,0545	9
		Simplicity	0,1456	3	0,0336	13
Facilitating Conditions	0,2215	Financial Resources	0,5366	1	0,1189	2
		Knowledge & Awareness	0,2072	3	0,0459	10
		Infrastructure	0,2561	2	0,0567	7
Social Influence	0,2274	Culture	0,4980	1	0,1132	3
		Government Support	0,2415	3	0,0549	8
		Sustainability Expectations	0,2605	2	0,0592	6
Perceived Risks	0,149	Overreliance	0,6125	1	0,0913	5
		Bias	0,2357	2	0,0351	12
		Hallucination	0,1519	3	0,0226	14

V. FINDINGS

Findings of the AHP analysis presented the main drivers and sub-drivers of generative AI acceptance in agriculture as given in Table 4. Prior to interpreting the results, the consistency of all pairwise comparison matrices were verified. All aggregate Consistency Ratio (CR) values were well below the 0.10 threshold required for acceptable consistency [27]. For the ChatGPT-generated judgements, the main driver CR was 0.0012, and sub-driver CRs ranged from 0.0014 to 0.0205. All individual farmer matrices also passed the CR threshold, with maximum individual CR values of 0.0488 across all matrices. These results confirm the internal consistency of the synthetic expert judgements used in the analysis. The most significant main driver is perceived ease of use with social influence ranked as second by synthetic farmers judgements. The least significant main driver is considered perceived risks and the fourth ranked is perceived usefulness. Perceived ease of use being the most significant driver is an expected result with the highly popularity of generative AIs and easier access to them through mobile systems. The social influence as another significant driver to acceptance in agriculture is also aligned with individuals' peers and social ecosystem being a significant factor in their decision-making. Especially in the context of agriculture, rural communities' practices are more linked with higher social ties. The low ranking of perceived usefulness and perceived risks indicates that the benefits and risks of generative AI tools are overshadowed by the lack of knowledge on generative AIs uses in agriculture. This result can be explained by the generative AI's uses and adoption being at the earlier stages in agriculture.

The risks and benefits being unknown to farmers present a research gap that highlights the need for awareness campaigns, consistent with the SLR finding that acceptance-focused studies remain scarce relative to theoretical and design studies.

The most significant sub-driver is accessibility, reflecting the growing adoption and easier access to generative AI across industries. Financial resources and culture ranked second and third, reflecting the role of economic constraints and traditional values in technology acceptance. The least significant sub-drivers were traceability, hallucination, and simplicity.

VI. CONCLUSION

The contribution of this study consists of systematically identifying the drivers of generative AI acceptance in agriculture, generating synthetic experts on agriculture and evaluating the proposed framework by synthetic farmer profiles judgements with AHP analysis. The findings of the AHP analysis present preliminary indications of potentially significant drivers to investigate in future empirical research on generative AI acceptance in agriculture. The proposed framework of drivers reveals an ease-of-use acceptance model rather than a risk based. The results provide insight to AI developers and policy creators to focus on practical uses and easier access on generative AI use as main priority. Another significant contribution is using generative AI to demonstrate the feasibility of synthetic experts for early-stage assessment of decision frameworks.

Limitations include the non-deterministic nature of LLM outputs, potential underrepresentation of agricultural workers from developing regions in training data, and the use of a single generative AI model. Future work will address these through multi-LLM benchmarking, human expert validation, and a survey-based framework assessment. Supplementary materials including synthetic farmer profiles and generation prompts are publicly available at https://doi.org/10.5281/zenodo.20044579.

References

[1] B. Kitchenham, "Procedures for performing systematic reviews," *Keele, UK, Keele University*, vol. 33, no. 2004, pp. 1–26, 2004.

[2] O. M. Awoleye *et al.*, "Leveraging Artificial Intelligence-Powered Virtual Assistant for Information Retrieval in Indigenous Agriculture: Insights from Nigeria," in *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Padua Italy: ACM, Jul. 2025, pp. 3095–3097.

[3] N. A. Akbar, B. Lenzitti, and D. Tegolo, "A Novel Approach for Leveraging Agent-Based Experts on Large Language Models to Enable Data Sharing Among Heterogeneous IoT Devices in Agriculture," in *AIxIA 2024 – Advances in Artificial Intelligence*, vol. 15450, A. Artale, G. Cortellessa, and M. Montali, Eds., Cham: Springer Nature Switzerland, 2025, pp. 12–22.

[4] B. Mehra and J. Anitha, "Empowering Indian Farmers with Multilingual AI: A Voice-Enabled ChatBot Using Dhenu2," in *2025 Second International Conference on Cognitive Robotics and Intelligent Systems (ICC-ROBINS)*, IEEE, 2025, pp. 556–560.

[5] E. E. Farmer, D. Brown, M. A. Gore, and H. A. Tufan, "Applying large language models to extract information from crop trait prioritization studies," *Plants People Planet*, p. ppp3.70075, Aug. 2025.

[6] K. Ramalakshmi, S. Shirly, R. Venkatesan, G. N. Sundar, and S. Kumar, "AI-based Bonsai Species Classification and Care Recommendation System using YOLO," in *2025 International Conference on Intelligent Computing and Control Systems (ICICCS)*, IEEE, 2025, pp. 642–647.

[7] P. D. Reddy, K. S. S. Reddy, P. Jayanth, B. P. Kakarla, and R. M. Balakrishnan, "Agri Assist: An AI Integrated Farmer Assistant," *Procedia Computer Science*, vol. 258, pp. 3510–3522, 2025.

[8] K. B. Suhas, A. S. Krishna, M. S. Roopesh, K. Teja, A. G. Kumar, and K. Nandanan, "Generative AI for community empowerment: transforming livelihood opportunities in a rural Indian village," in *2024 Fourth International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT)*, IEEE, 2024, pp. 1–6.

[9] H. Nikam *et al.*, "Enhancing Agricultural Practices with AgroDialogue: A Chatbot Framework," in *2024 IEEE Pune Section International Conference (PuneCon)*, IEEE, 2024, pp. 1–11.

[10] K. Herrera, J. Miranda, and D. Mauricio, "Milchbot: app to support the process of feeding and caring for dairy cows in Peru," *AGRIS on-line Papers in Economics and Informatics*, vol. 14, no. 4, pp. 27–37, 2022.

[11] M. V. Leite, J. M. Abe, M. L. H. Souza, and I. de Alencar Nääs, "Enhancing Environmental Control in Broiler Production: Retrieval-Augmented Generation for Improved Decision-Making with Large Language Models," *AgriEngineering*, vol. 7, no. 1, p. 12, 2025.

[12] F. Makudza, P. Dangaiso, S. Muparangi, E. Mtisi, N. Makiwa, and G. Makandwa, "Leveraging climate resilience through AI chatbot farming adoption among rural Zimbabwean women," *Cogent Social Sciences*, vol. 11, no. 1, p. 2546520, Dec. 2025.

[13] I. Ajzen, "The theory of planned behavior," *Organizational behavior and human decision processes*, vol. 50, no. 2, pp. 179–211, 1991.

[14] F. D. Davis, "Perceived Usefulness, Perceived Ease of Use and User Acceptance of Information Technology," *MIS quarterly*, 1989.

[15] E. M. Rogers, *Diffusion of innovations*, 5th Edition. New York: The Free Press, 2003.

[16] L. Tornatzky and M. Fleischer, *The Process of Technology Innovation*. MA: Lexington, 1990.

[17] V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis, "User acceptance of information technology: Toward a unified view," *MIS quarterly*, pp. 425–478, 2003.

[18] M. T. Çaldağ, "AI Pair Programming Acceptance: A Value-Based Approach with AHP Analysis," in *2024 10th International Conference on Control, Decision and Information Technologies (CoDIT)*, IEEE, 2024, pp. 556–561.

[19] Z. Ji *et al.*, "Survey of Hallucination in Natural Language Generation," *ACM Comput. Surv.*, vol. 55, no. 12, p. 248:1-248:38, Mar. 2023.

[20] L. P. Argyle, E. C. Busby, N. Fulda, J. R. Gubler, C. Rytting, and D. Wingate, "Out of one, many: Using language models to simulate human samples," *Political Analysis*, vol. 31, no. 3, pp. 337–351, 2023.

[21] M. T. Çaldağ, "Challenges on Artificial Expert Acceptance in AHP Analysis," presented at the 2025 11th International Conference on Control, Decision and Information Technologies (CoDIT), IEEE, 2025, pp. 1007–1012.

[22] M. Sarstedt, S. J. Adler, L. Rau, and B. Schmitt, "Using large language models to generate silicon samples in consumer and marketing research: Challenges, opportunities, and guidelines," *Psychology & Marketing*, vol. 41, no. 6, pp. 1254–1270, 2024.

[23] J. Bisbee, J. D. Clinton, C. Dorff, B. Kenkel, and J. M. Larson, "Synthetic replacements for human survey data? The perils of large language models," *Political Analysis*, vol. 32, no. 4, pp. 401–416, 2024.

[24] A. Li, H. Chen, H. Namkoong, and T. Peng, "Llm generated persona is a promise with a catch," *arXiv preprint arXiv:2503.16527*, 2025.

[25] L. Salewski, S. Alaniz, I. Rio-Torto, E. Schulz, and Z. Akata, "In-context impersonation reveals large language models' strengths and biases," *Advances in neural information processing systems*, vol. 36, pp. 72044–72057, 2023.

[26] S. Chapala, M. Mironov, and S. Deng, "Mitigating Social Desirability Bias in Random Silicon Sampling," *arXiv preprint arXiv:2512.22725*, 2025.

[27] T. L. Saaty, "How to make a decision: the analytic hierarchy process," *European journal of operational research*, vol. 48, no. 1, pp. 9–26, 1990.