AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology

Uday Surendra Yandamuri

Independent Researcher, yudaysurendra@gmail.com, 0009-0003-8655-9322

Abstract
The hospitality technology industry has advanced infrastructure with integrated information systems. Well-designed AI modules—within an established Data Ecosystem and Development Framework—can support various organizational roles and enhance efficiency, speed, and reliability. Managerial decisions often require the help of data- and/or model-driven Decision Support Systems (DSS), such as reporting dashboards, staff scheduling tools, inventory management systems, and revenue management solutions. Demand forecasting, operational-level staffing, inventory management, dynamic pricing, and service-level optimization are primary decision areas. However, several limits impede DSS effectiveness: lack of user trust, excessive cognitive load, inadequate validation and monitoring, and poorly executed data-driven embeddings. Well-designed AI DSS addresses the challenges.
The proposed DSS concepts combine Decision Science fundamentals with specific Hospitality Technology requirements. Five central dimensions apply: key phases in substantial decision-making processes; different Decision Support levels in the operational-executive domain; the roles of decision-makers—receivers, planners, and validators; Human-in-the-loop methodology; and theoretical pillars, including Robustness, Interpretability, and Explainability. With a suitable architecture, Data Ecosystem, specific methods, and Human-AI collaboration, an AI DSS can assist operational optimization.
Keywords : Artificial Intelligence; Hospitality; Decision Support Systems; Operational Optimization; Demand Forecasting; Staff-Scheduling; Inventory Management; Dynamic Pricing; Service Levels; Transfer Learning

1. Introduction

Recent COVID-19 waves have dented consumer confidence and demand in the hospitality sector, further exacerbating existing vulnerabilities and uncertainties. Although demand may be only temporarily affected, it differs significantly from pre-pandemic times: travelers are now more price-sensitive and customer preferences can shift rapidly. At the same time, operational costs remain high due to staffing shortages and inflation. Hospitality companies need much more accurate short-term demand forecasting to cope with these challenges while embracing new technologies.

In the context of hospitality technology, decision support systems (DSSs) powered by artificial intelligence (AI) have strong potential to make a difference. AI, combined with the rich troves of historical, geolocation, and social media data generated in the sector, provides the necessary ingredients to improve operational decision making in digital ecosystems. The approaches currently being implemented or developed in the area of decision support systems cover key hospitality functions such as demand forecasting, staff scheduling, inventory management, dynamic pricing, and service level.The information integration challenge follows a schema-on-read approach, allowing for separate designs of the multiple hospitality technology components while ensuring proper aligned synthesis over time. The quality of feature engineering dictates the success of AI efforts. Features that capture temporal evolution, information propagation delays, and event-delivered information are instrumental in achieving higher data richness.

Fig 1: AI Agent in Hospitality

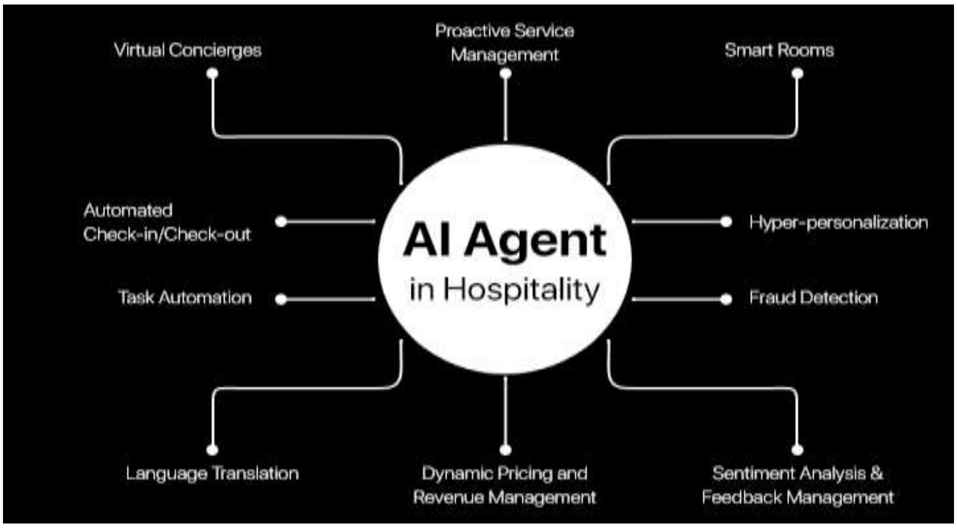

2. Theoretical Foundations of Decision Support Systems in Hospitality

Decision support systems (DSS) enable human decision-makers to make informed choices through data-driven insights. DSS can be classified into model-driven, data-driven, or hybrid systems. Interwoven with decision science, a discipline devoted to understanding how decisions are made, a DSS architecture enables the deployment of various AI methods to assist specific human decision processes or operations. Four primary operations within hospitality management demand decision support: forecasting, scheduling, inventory control, and pricing. Expert knowledge in these areas is expressed through theory and models of operational research, queueing systems, and revenue management.

Operational-level forecasting supports demand predictions, the basis for staffing, inventory, and service-level decisions. Human planners assign schedules that utilize labor to profitably meet expected demand and its variability. Inventory decisions determine optimal stock levels of key assets that either directly fulfill demand (rooms, tables) or are prerequisites for enabling demand completion (food and beverage). Responding to fluctuating leisure and business demand patterns, dynamic pricing leverages the gap between price and variable cost to maximize occupancy and revenue.

Equation 1: Forecasting: core error metrics (explicitly referenced in the paper)

Assume:

- Actual demand at time t: y_t
- Forecast: $\hat{y}_t$
- Number of evaluated points: T

1.1 Mean Absolute Error (MAE)

Step-by-step

1. Forecast error:

$$e_t = y_t - \hat{y}_t$$

2. Absolute error:

$$|e_t| = |y_t - \hat{y}_t|$$

3. Average absolute error over T points:

$$\text{MAE} = \frac{1}{T}\sum_{t=1}^{T}|y_t - \hat{y}_t|$$

1.2 Root Mean Squared Error (RMSE)

Step-by-step

1. Squared error:

$$e_t^2 = (y_t - \hat{y}_t)^2$$

2. Mean squared error (MSE):

$$\mathrm{MSE} = \frac{1}{T}\sum_{t=1}^{T}(y_t - \hat{y}_t)^2$$

3. Root of MSE:

$$\mathrm{RMSE} = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(y_t - \hat{y}_t)^2}$$

Visuals included: demand forecast plot + MAE/RMSE bar chart (synthetic example) in the downloads above.

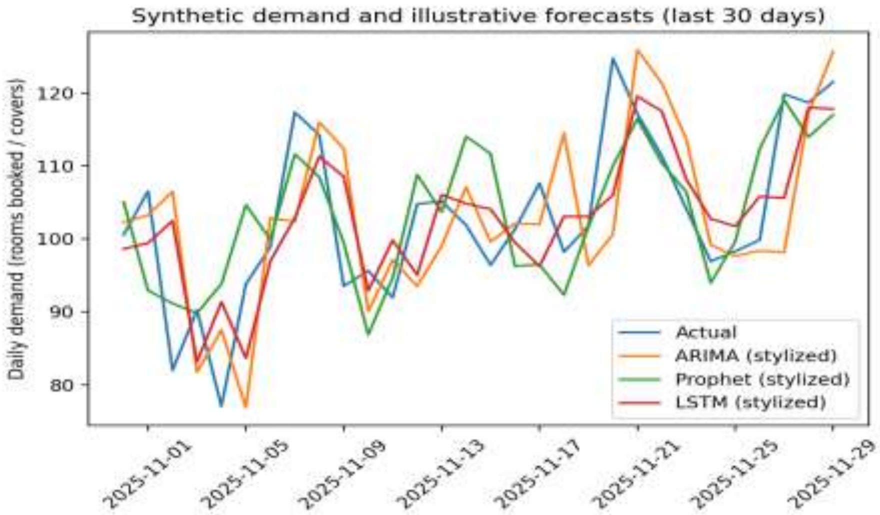

3. Architecture of AI-Driven Decision Support Systems

An AI-driven decision support system (DSS) comprises six logical layers: data ingestion, processing, AI model serving, inference, and presentation. Interoperability across layers ensures a seamless flow of data, enabling all subsystems to collaborate closely throughout the operational cycle. Data governance, privacy, and security considerations guide the deployment of organization-wide DSS features.

DSS requires a feature store (or engine) for the systematic creation, curation, storage, and management of task-specific features. Major components that provide a complete AI DSS ecosystem include a data lake/warehouse, feature store, model registry, orchestration engine, user interface, and application programming interfaces (APIs) for external integration, data-sharing governance, and systematic monitoring. A well-defined data lineage ensures the analytical capabilities are scalable, while an explainable artificial intelligence approach helps bridge the gap between data-driven and human-centric decision-making processes in hospitality operations.

A robust Decision Support System (DSS) for hospitality operations relies on an integrated AI ecosystem that enables reliable, scalable, and transparent decision-making. Central to this ecosystem is a feature store or feature engine, which supports the systematic creation, curation, storage, and reuse of task-specific features to ensure consistency between model training and deployment. Complementing this are foundational components such as a data lake or data warehouse for unified data storage, a model registry for versioning and governance of analytical models, and an orchestration engine to automate data pipelines and model workflows. User interfaces and application programming interfaces (APIs) facilitate seamless interaction, external system integration, and secure data sharing across stakeholders. Strong data governance mechanisms, including end-to-end data lineage and continuous monitoring, ensure analytical processes remain auditable, scalable, and trustworthy. Furthermore, embedding explainable artificial intelligence (XAI) within the DSS bridges the gap between data-driven insights and human-centric decision-making, enhancing transparency, accountability, and adoption of AI-driven recommendations in hospitality management.

Fig 2: Architecting an AI-Driven Decision Support System

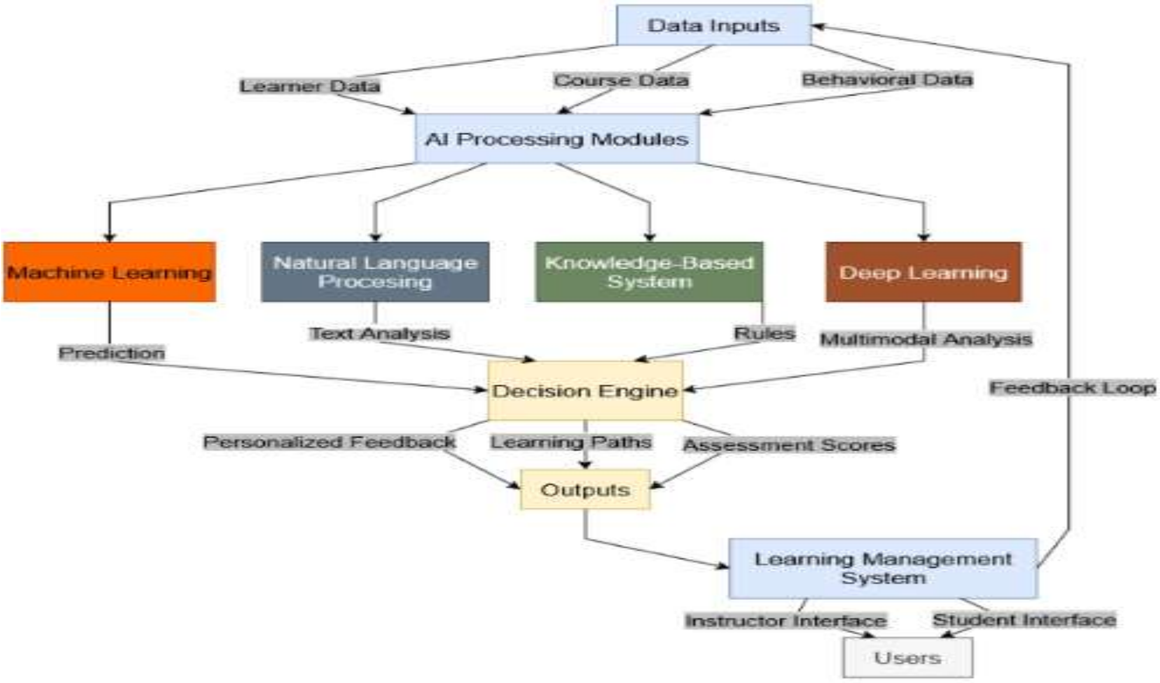

4. Data Ecosystem and Feature Engineering in Hospitality

A data ecosystem supports decision-making in hospitality operations through a collection of data sources that progressively gathered information on different aspects of hospitality operations. Important data sources include property management systems (PMS), point-of-sale systems (POS), central reservation systems (CRS), revenue management systems (RMS), Internet of Things (IoT) sensor data, and online customer reviews. Every aspect of the decision-making process is supported by multiple data sources. This involves naturally spiraling a large amount of information from different sources into a common privacy-preserving data architecture, ensuring high data quality while controlling associated costs, and growing the hospitality technology stack toward higher additional data quality and predictive performance.

Feature engineering enables operationalization of the data foundations through big data pipelines, responsible for feature synthesis and validation, as well as data integration into a common representation format. Integrating all available information into a single representation format along its temporal dimension with the appropriate provenance supports time- and event-based model training.

Equation 2: Time-series models mentioned: ARIMA / SARIMA / Prophet (standard forms)

Let B be the backshift operator: $By_t = y_{t-1}$.

Step A — Differencing (to make series stationary)

If d = 1:

$$\nabla y_t = (1 - B)y_t = y_t - y_{t-1}$$

If d times:

$$\nabla^d y_t = (1 - B)^d y_t$$

Define $x_t = \nabla^d y_t$.

Step B — AR and MA polynomials

AR (p):

$$\phi(B)x_t = \left(1 - \phi_1 B - \cdots - \phi_p B^p\right)x_t$$

MA (q):

$$\theta(B)\varepsilon_t = \left(1 + \theta_1 B + \cdots + \theta_q B^q\right)\varepsilon_t$$

Step C — Combine into ARIMA

$$\phi(B)\, x_t = \theta(B)\, \varepsilon_t$$

i.e.

$$\phi(B)\, (1 - B)^d y_t = \theta(B)\, \varepsilon_t$$

5. AI Methods for Operational Optimization

The AI portfolio for operational optimization encompasses a diverse set of methods and tasks, including demand forecasting, staff scheduling and allocation, food and beverage inventory management, anomaly detection, dynamic pricing, service-level optimization, and other supporting decision-making activities. The core methods comprise time-series forecasting models (ARIMA, SARIMA, Prophet, LSTM), GBM and neural-network-based models for non-DT-structured data sources (e.g., predictors from hotel guest-generated reviews), multi-agent deep reinforcement learning (DRL) for revenue management, sophisticated Optimizers such as open-source libraries (OptaPlanner, Google OR-Tools) or commercial engines (FICO, Kyndryl, PROS) for operational research and crew-crew assignment problems, ad hoc algorithms for F&B inventory-level forecasting based on historical POS data inputs, and machine learning (ML) as well as statistical baseline anomaly detection models.

Each method has been linked with specific operational goals: accurate demand forecasts contribute to inventory-planning accuracy (e.g., forecast-driven F&B inventory replenishment and F&B stock-out mitigation), improved staff-scheduling quality and staff-utilization efficiency (capacity losses and service-level degradation) reduce operational costs and improve profitability, and accurate dynamic-pricing models increase hotel-room revenue. In addition to serving as a foundation for building actionable tasks, the underlying AI-driven DSS offers essential supporting decision-support capabilities to elevate the quality of DM process execution and both short- and long-term decision-making.

Fig 3:

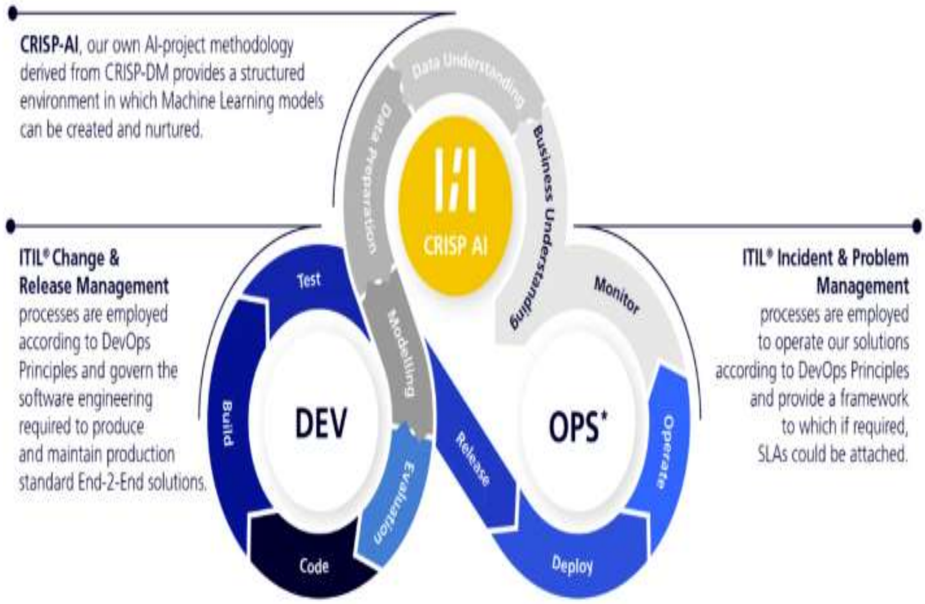

6. Human–AI Collaboration in Hospitality Operations

Integrating human-and-AI collaboration into operational optimization determines the essential roles of decision-makers and highlights the importance of trust, user interfaces, and governance of human overrides. Supporting hospitality operations with AI methods is still a work in progress, as empirical studies cover only part of the service chain and deploy only a subset of available models. The proposed solutions so far incorporate decision support as a design and evaluation criterion, but the actual interfaces and interaction with final users have received limited attention. User interfaces, determinative for successful implementation and adoption, require adequate explainability, control for reduced cognitive load, and integration with other information systems.

Moreover, controlling for the human-in-the-loop aspect in experimental design and model development is essential but difficult, due to the different speeds of technology development and adoption. Therefore, governance processes for human overrides on suggested decisions and their consequent feedback on model retraining remain critical open questions. Candidate processes rely on validation, monitoring, and feedback loops aligned with operational decision-making routines. Well-designed validation processes enable the continuous adjustment of margins required to maintain service levels or business profitability, accommodating the final user experiences and opinions in the prediction and suggestion architecture. Monitoring resists the decision-support model drift and quantifies when a scheduled maintenance task is required. Feedback loops close the circuit of the decision-support perceptual system, allowing a continuous improvement process by integrating users' decisions and controls back into ML refinements.

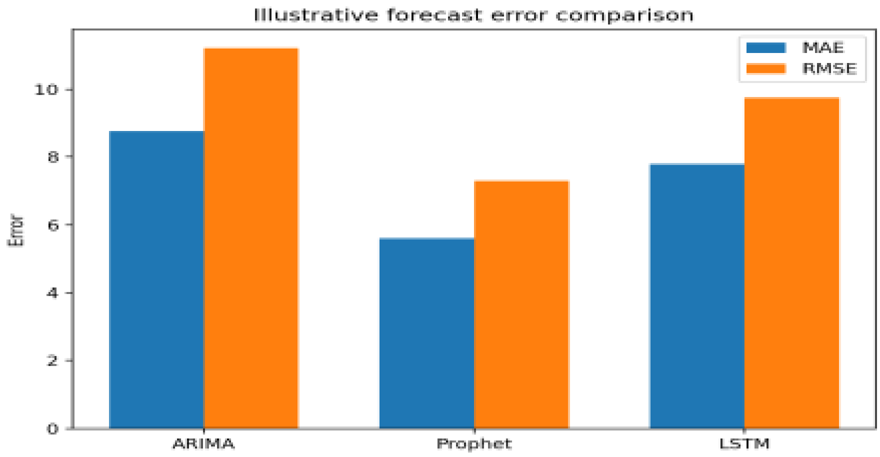

7. Evaluation Metrics and Experimental Design

AI DSS aim to enhance hospitality operations by integrating decision support within business processes, with accuracy and operational goals as primary guiding principles. A diverse portfolio of quantitative metrics ensures that objectives are met across different tasks and methods, while multiple experimental designs enable rigorous technical validation.

Evaluation metrics follow operational goals: forecast accuracy influences demand forecasting; hospitality profitability and guest satisfaction inform demand and service-level models; and labor utilization and staffing shortages address employee scheduling. Quantitative time-series forecasting is measured by mean absolute error, root mean squared error, and out-of-sample prediction precision, with external benchmark performance established as an accepted validation approach. A/B testing evaluates frontal models for hospitality profitability (e.g., revenue management) and service-level metrics (e.g., guest satisfaction and waiting times). These hospitality-specific metrics are synthesized in a target-occupancy surface model to determine the optimal dual occupancy and average construction for real-time forecasting. The overall accuracy of staff scheduling is represented by labor utilization, complemented by the absence of staffing shortages.

Forward-looking methods (e.g., inventory management, personalized marketing, and pricing decision strategy) draw on parameters generated by other AI DSS. Backtests for reinforcement-learning models incorporate choice accuracy of candidate agents and profitability for personalized-marketing campaigns. Scalability in scheduling for demand pattern anomalies is assessed by a sequential var-ma model integrated with a gradient-boosting labor-scheduling AI. Robustness testing assesses performance deviation with out-of-control inputs in demand-patterns, language sentiment, and queuewaiting-time prediction. Transfer-ability of the models is assessed by spatial-temporal k-fold cross-validation, while templates ensure reproducibility in application to other hospitality operations.

Fig 4: Evaluation Metrics and Experimental Design

8. Implementation Challenges and Risk Management

The successful deployment of AI action-oriented decision support systems in hospitality depends on careful consideration of eight issues: the availability and quality of data, compliance with privacy laws, protection against data breaches, strategies for educating and supporting partners involved in the provision of physical services, the management of system-generated recommendations, the integration of the system within a hospitality group's information technology stack, the management of vendor selection processes, and the risk of scalability limitations. These concerns are explored below.

Efficient transfer learning from an established AI action-oriented decision support system to a new hotel is subject to proper data governance that tracks the origin of each piece of information, ensures compliance with data use agreements, and safeguards identity and emotional privacy through data anonymization. Data governance must also guard against inappropriate use of data. AI action-oriented decision support systems can be manipulated to produce unacceptable service decisions — for example, selecting poisonous ingredients for food or setting service quotas that exceed appropriate levels for guests or destinations — if decisions are delayed or even suppressed. Adequate training, continuous support, and judicious monitoring of user decisions can help mitigate such risks, which are usually perceived as minimal or negligible.

Equation 3: LSTM forecasting (step-by-step gate equations)

(1) Forget gate

$$f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f)$$

(2) Input gate

$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i)$$

(3) Candidate cell content

$$\tilde{c}_t = \tanh(W_c x_t + U_c h_{t-1} + b_c)$$

3.2 Cell state update

Step-by-step

1. Keep some old memory: $f_t \odot c_{t-1}$
2. Add some new memory: $i_t \odot \tilde{c}_t$
3. Combine:

$$c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t$$

3.3 Output gate and hidden state

$o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o)$ $h_t = o_t \odot \tanh(c_t)$

Forecast is then typically:

$$\hat{y}_{t+1} = W_y h_t + b_y$$

9. Conclusion

Hospitality decision makers are increasingly recognising the potential of AI-driven DSS and seeking to adopt them. Embedding these tools within a robust decision-making process can significantly alleviate reservations and concerns regarding their potential impacts. However, even when well designed, methods typically covering only a single area of operations—such as forecasting, inventory management, or revenue—may lead to local optimisations that fail to consider the entire business. The strength of such a decision-support architecture lies, therefore, in its holistic nature, encompassing multiple aspects of hospitality and integrating specialist management knowledge with AI capabilities. By allowing domain experts to select, train, tune, and monitor a wide range of methods covering multiple aspects of hospitality operations, the approach enhances robustness and widens the breadth of application.

Early-stage experimental results from implementations in large European hotel groups demonstrate the significant benefits of this AI, human-in-the-loop, decision-support approach implemented across multiple operational areas. A backtesting approach has confirmed the importance of demand forecasting and dynamic inventory management and illustrated the revenue impacts of a simple, hierarchical reinforcement-learning agent for service-level forecasting. Given the variable cost of hospitality, staff scheduling remains an operational area where decisive action and human expertise can make a difference to the bottom line. Nevertheless, sophisticated dashboard systems are already showing greater capability than operational managers, and the search for better methods remains essential. Continued human involvement in the application of the AI tools, supported by a process of good governance, is crucial to maximising profits.

10. References

[1] Al Shehhi, A., & Karathanasopoulos, A. (2020). Forecasting hotel prices (ADR) using statistical and artificial intelligence methods: Evidence from Persian Gulf cities. Tourism Economics, 26(8), 1279–1301.
[2] Nagabhyru, K. C. (2024). Data Engineering in the Age of Large Language Models: Transforming Data Access, Curation, and Enterprise Interpretation. Computer Fraud and Security.
[3] Borghi, M., & Mariani, M. M. (2023). The impact of service robots on customer satisfaction: Evidence from hospitality services. Psychology & Marketing, 40(11), 2370–2386.
[4] Borghi, M., Mariani, M. M., & Wirtz, J. (2024). Asymmetrical influences of service robots' perceived intelligence and warmth on consumer responses in hospitality. Journal of Travel Research, 63(x), xx–xx.
[5] Aitha, A. R. (2024). Generative AI-Powered Fraud Detection in Workers' Compensation: A DevOps-Based Multi-Cloud Architecture Leveraging, Deep Learning, and Explainable AI. Deep Learning, and Explainable AI (July 26, 2024).
[6] Chen, S., Ngai, E. W. T., Ku, Y., Xu, Z., Gou, X., & Zhang, C. (2023). Prediction of hotel booking cancellations: Integration of machine learning and probability model based on interpretable feature interaction. Decision Support Systems, 170, 113959.
[7] Contessi, D., Viverit, L., Pereira, L. N., & Heo, C. Y. (2024). Decoding the future: Proposing an interpretable machine learning model for hotel occupancy forecasting using principal component analysis. International Journal of Hospitality Management, 119, 103802.
[8 Keerthi Amistapuram. (2024). Federated Learning for Cross-Carrier Insurance Fraud Detection: Secure Multi-Institutional Collaboration. Journal of Computational Analysis and Applications (JoCAAA), 33(08), 6727–6738. Retrieved from https://www.eudoxuspress.com/index.php/pub/article/view/3934.
[9] Gajić, T., Blešić, I., Petrović, M. D., & Vukolić, D. (2024). Innovative approaches in hotel management: Integrating artificial intelligence (AI) and the internet of things (IoT) to enhance operational efficiency and sustainability. Sustainability, 16(17), 7279.
[10] Henriques, H., & Pereira, L. N. (2024). Hotel demand forecasting models and methods using artificial intelligence: A systematic literature review. Tourism & Management Studies, 20(x), xx–xx.
[12 Deep Learning-Driven Optimization of ISO 20022 Protocol Stacks for Secure Cross-Border Messaging. (2024). MSW Management Journal, 34(2), 1545-1554.
[13] Infante, A., García, M., & Ferrer, J. (2021). Adoption of AI-enabled self-service technologies in tourism and hospitality: A review and research agenda. Journal of Hospitality and Tourism Technology, 12(x), xx–xx.
[14] Segireddy, A. R. (2024). Machine Learning-Driven Anomaly Detection in CI/CD Pipelines for Financial Applications. Journal of Computational Analysis and Applications, 33(8).
[15] Kaya, K., Ergen, B., & Delen, D. (2022). Demand forecasting model using hotel clustering findings and attention-based deep learning. Information Processing & Management, 59(3), 102816.
[16] Knani, M., Ben-Hamadou, A., & Guesmi, K. (2022). Artificial intelligence in tourism and hospitality: A bibliometric analysis and future research directions. Tourism Management Perspectives, 44, 101012.
[17] Bachhav, P. J., Suura, S. R., Chava, K., Bhat, A. K., Narasareddy, V., Goma, T., & Tripathi, M. A. (2024, November). Cyber Laws and Social Media Regulation Using Machine Learning to Tackle Fake News and Hate Speech. In International Conference on Applied Technologies (pp. 108-120). Cham: Springer Nature Switzerland.
[18] Mariani, M. M., & Borghi, M. (2023). Artificial intelligence and analytics in hospitality and tourism: A critical reflection and future directions. International Journal of Contemporary Hospitality Management, 35(x), xx–xx.
[19] Kannan, S., & Saradhi, K. S. Generative AI in Technical Support Systems: Enhancing Problem Resolution Efficiency Through AIDriven Learning and Adaptation Models.
[20] Park, E., Kim, S., & Lee, J. (2024). Service robots in hospitality and tourism research: A bibliometric analysis and research agenda. SAGE Open, 14(x), 21582440241258281.
[21] Rodrigues, M., Miguéis, V., Freitas, S., & Machado, T. (2023). Machine learning models for short-term demand forecasting in food catering services: A solution to reduce food waste. Journal of Cleaner Production, 413, 140265.
[22] Chakilam, C., Suura, S. R., Koppolu, H. K. R., & Recharla, M. (2022). From Data to Cure: Leveraging Artificial Intelligence and Big Data Analytics in Accelerating Disease Research and Treatment Development. Journal of Survey in Fisheries Sciences. https://doi.org/10.53555/sfs.v9i3.3619.
[23] Saitta, S., D'Amico, V., & Farinella, G. (2024). Dynamic price prediction for revenue management systems in the hospitality sector. In Proceedings of the 13th International Conference on Data Science, Technology and Applications (DATA 2024) (pp. 218–228). SCITEPRESS.
[24] Samala, N., Katkam, B. S., Bellamkonda, R. S., & Rodriguez, R. V. (2022). Impact of AI and robotics on the tourism and hospitality industry: A systematic literature review. Journal of Hospitality and Tourism Technology, 13(x), xx–xx.
[25] Challa, S. R. (2024). The Future of Banking and Lending: Assessing the Impact of Digital Banking on Consumer Financial Behavior and Economic Inclusion. Available at SSRN 5151025.
[26] Soori, M., Karimi Ghaleh Jough, F., Dastres, R., & Arezoo, B. (2024). AI-based decision support systems in Industry 4.0: A review. Journal of Economy and Technology, 1(x), xx–xx.
[27] Tuomi, A., Tussyadiah, I., & Stienmetz, J. (2021). Service robots and the changing hospitality workplace: Implications for operations and employee outcomes. International Journal of Hospitality Management, 95, 102912.
[28] Pamisetty, V. (2024). AI-Driven Decision Support for Taxation and Unclaimed Property Management: Enhancing Efficiency through Big Data and Cloud Integration. European Journal of Analytics and Artificial Intelligence (EJAAI) p-ISSN 3050-9556 en e-ISSN 3050-9564, 2(1).

[29] Wiastutti, R. D., Ivanov, S., & Webster, C. (2024). Service robots in the hotel industry: Customer acceptance across hotel brand contexts. Tourism and Hospitality Management, 30(3), 375–388.
[30] A Scalable Web Platform for AI-AugmentedSoftware Deployment in Automotive Edge Devices via Cloud Services. (2024). American Advanced Journal for Emerging Disciplinaries (AAJED) ISSN: 3067-4190, 2(1).
[31] Xu, J., Hsiao, A., Reid, S., & Ma, E. (2023). Working with service robots? A systematic literature review of hospitality employees' perspectives. International Journal of Hospitality Management, 112, 103523.
[32] Guntupalli, R. (2024). Enhancing Cloud Security with AI: A Deep Learning Approach to Identify and Prevent Cyberattacks in Multi-Tenant Environments. Available at SSRN 5329132.
[33] Zhang, Y., Liu, X., & Li, J. (2023). Leveraging online reviews for hotel demand forecasting: A deep learning approach. Information Processing & Management, 60(5), 103192.
[34] Rongali, S. K. (2024). Federated and Generative AI Models for Secure, Cross-Institutional Healthcare Data Interoperability. Journal of Neonatal Surgery, 13(1), 1683–1694.