



Intelligent Tenant Sales Auditing and Anomaly Detection for Turnover-Based Lease Management in Saudi Retail Real Estate

Abdullah Alijefri^{1*}

¹Cenomi Centres – Manager, Lease Contracts Management, Riyadh, Saudi Arabia

*Corresponding Author

Abdullah Alijefri

Cenomi Centres – Manager,
Lease Contracts Management,
Riyadh, Saudi Arabia

Article History

Received: 03.08.2026

Accepted: 01.10.2026

Published: 08.10.2026

Abstract: The expansion of modern retail destinations in Saudi Arabia has increased the complexity of tenant and lease administration, particularly where rental arrangements include turnover-based components linked to tenant sales. Conventional tenant sales auditing depends heavily on periodic declarations, manual reconciliation, sampling, and retrospective verification. These practices can be difficult to scale across multi-mall portfolios and may delay the identification of abnormal sales patterns. This study proposes an intelligent tenant sales auditing framework that combines machine learning-based anomaly detection, retail sales forecasting, tenant benchmarking, contractual data analysis, and risk-based audit prioritisation. The proposed framework integrates tenant point-of-sale information, declared sales, historical performance, footfall indicators, lease conditions, payment records, and tenant-category benchmarks. A simulated evaluation demonstrates how the framework can classify tenants according to anomaly risk and support targeted audit intervention. The study does not claim that simulated results represent actual Saudi mall performance; they illustrate the analytical logic and potential operational value of the proposed approach. The framework contributes to research at the intersection of artificial intelligence, retail analytics, and commercial real estate management by shifting tenant sales auditing from periodic and reactive verification toward continuous, data-driven monitoring. The proposed model can support Saudi retail property operators in improving sales transparency, strengthening lease governance, prioritising audit resources, and developing more intelligent turnover-based lease management practices.

Keywords: Artificial Intelligence, Tenant Sales Auditing, Anomaly Detection, Turnover-Based Lease, Retail Real Estate, Saudi Arabia, Machine Learning, Predictive Analytics, Lease Management, Retail Analytics, Commercial Property Management.

Copyright © 2026 The Author(s): This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY-NC 4.0) which permits unrestricted use, distribution, and reproduction in any medium for non-commercial use provided the original author and source are credited.

1. INTRODUCTION

Saudi Arabia's retail real estate sector is undergoing rapid structural change as economic diversification, urban development, digital commerce, and experience-led destinations reshape the relationship between consumers, retailers, and

property owners. Recent market evidence indicates that consumer spending reached approximately SAR 1.41 trillion in 2024, while point-of-sale transactions reached SAR 668 billion, demonstrating the scale of the commercial ecosystem in which tenant sales are generated and monitored (Knight Frank, 2025).

Citation: Alijefri, A. (2026). Intelligent Tenant Sales Auditing and Anomaly Detection for Turnover-Based Lease Management in Saudi Retail Real Estate; *Glob Acad J Econ Buss*, 8(5), 1641-1651. <https://doi.org/10.36348/gajeb.2026.v08i05.051>

Large shopping centres increasingly combine conventional retail with food and beverage, entertainment, hospitality, leisure, and experiential uses. This diversification increases the commercial complexity of lease administration and creates a stronger need for accurate, timely, and auditable tenant performance information.

Turnover-based leases are particularly relevant because they link some portion of occupancy cost to tenant sales. The mechanism can align landlord and tenant interests, yet it also creates information asymmetry. Retailers normally possess the most detailed transaction information, while landlords depend on contractual reporting, point-of-sale interfaces, declarations, financial statements, or audit rights to establish the sales base used for rent calculation. Any gap between actual commercial activity and information available to the property owner can create an audit risk.

Tenant sales auditing therefore has a dual purpose. It is not simply a financial control procedure; it is also a mechanism for validating lease compliance, identifying data inconsistencies, supporting commercial negotiations, and protecting the integrity of portfolio-level performance reporting. In a small property, a leasing team may investigate unusual results manually. In a multi-mall portfolio, however, the number of tenants, transaction records, lease clauses, reporting periods, product categories, and seasonal effects can make exhaustive manual review impractical. The challenge becomes one of prioritisation: management needs to identify which tenants require attention and why.

Artificial intelligence and machine learning offer a potential response. Retail research shows that machine learning can provide value when large and heterogeneous datasets must be analysed across many stores or products (Ma & Fildes, 2021; Grewal *et al.*, 2021). Retail forecasting research also demonstrates that sales series contain complex seasonality, volatility, promotions, and other external influences that make simple universal forecasting rules inadequate (Fildes *et al.*, 2022). These characteristics are directly relevant to tenant sales auditing because a legitimate fall in sales should not automatically be interpreted as a reporting anomaly. An intelligent auditing system therefore needs to establish an expected performance range before classifying deviations.

Anomaly detection provides a second important capability. Financial anomaly research has shown that unsupervised and semi-supervised techniques can be valuable where confirmed abnormal cases are rare and difficult to label comprehensively (Hilal *et al.*, 2022). In tenant

auditing, confirmed cases of misreporting or material discrepancy may be much smaller than the total number of ordinary tenant-month observations. This makes purely supervised classification restrictive. A combination of forecasting, peer benchmarking, and unsupervised anomaly detection can instead identify observations that depart from normal patterns.

The Saudi context adds another dimension. Retail property operators are increasingly exposed to digital transaction ecosystems, omni-channel consumer behaviour, and rapid development of new retail destinations. The continuing digitalisation of Saudi retail means that more operational data can potentially be used for intelligent audit controls (Knight Frank, 2025). At the same time, recent real-estate AI literature identifies persistent barriers involving data quality, model transparency, organisational capability, and governance. These challenges mean that an AI auditing system should not be treated as an autonomous enforcement mechanism. It should be designed as decision support with explainable alerts and human verification.

Despite growing research on AI in retail analytics, anomaly detection, and real-estate technology, comparatively limited work has examined the specific problem of intelligent tenant sales auditing for turnover-based lease management. Existing research commonly concentrates on consumer prediction, product demand forecasting, property valuation, or financial fraud detection. The connection between these capabilities and landlord-side tenant sales auditing remains underdeveloped. This study addresses that gap by proposing a conceptual framework tailored to Saudi retail real estate.

The research has five objectives: to identify the principal data and process factors influencing tenant sales auditing; to develop an integrated framework combining sales forecasting, anomaly detection, tenant benchmarking, and lease information; to establish a risk-scoring approach for prioritising tenant audits; to demonstrate the framework using transparent simulated data rather than fabricated commercial data; and to discuss implementation, governance, explainability, and future validation requirements.

The study asks four research questions: RQ1, which tenant-level indicators are most useful for identifying unusual sales behaviour? RQ2, how can forecasting and anomaly detection be combined to distinguish normal commercial volatility from potential reporting anomalies? RQ3, how can AI-generated risk scores improve audit prioritisation? RQ4, what organisational and governance conditions

are required for responsible adoption in Saudi retail real estate?

The paper contributes a structured AI-assisted auditing architecture, an integrated risk-scoring concept, and a practical implementation pathway. Its central proposition is that intelligent tenant auditing should not replace professional audit judgement; instead, it should make that judgement more targeted, timely, and evidence-based.

2. LITERATURE REVIEW

The literature relevant to intelligent tenant sales auditing spans retail sales forecasting, anomaly detection, artificial intelligence in retailing, and digital transformation in real estate. Together, these streams establish the technical foundation for an intelligent audit system while also revealing a research gap.

Retail sales forecasting is fundamental because an auditor cannot determine whether a tenant's reported sales are unusual without an estimate of normal performance. Ma and Fildes (2021) demonstrate that retail sales forecasting can benefit from meta-learning approaches that combine multiple forecasting methods and learn characteristics from large numbers of sales series. Their work is important for tenant auditing because tenants differ in category, scale, location, seasonality, and historical volatility. Fildes, Kolassa, and Ma (2022) similarly emphasise structural changes, instability, omni-retailing, and machine-learning adoption in retail forecasting. These findings support tenant-specific or peer-group forecasting.

Machine learning research in retailing provides a second foundation. Grewal *et al.*, (2021) discuss machine learning analytics and metrics in retail research and highlight the ability of machine learning to analyse large datasets while revealing patterns that traditional approaches may overlook. Policarpo *et al.*, (2021) review machine learning in e-commerce and identify applications associated with profit growth, customer analytics, and fraud detection. Although these studies are not specifically concerned with lease auditing, they establish that retail transaction environments provide suitable conditions for data-driven analytical systems.

Anomaly detection is particularly relevant where confirmed abnormal observations are limited. Hilal, Gadsden, and Yawney (2022) review anomaly detection techniques used in financial fraud and highlight the importance of semi-supervised and unsupervised approaches. This is conceptually appropriate for tenant sales auditing because most tenant-month records are expected to be legitimate,

while confirmed cases of material discrepancy are relatively infrequent.

Li and Jung (2023) review deep-learning methods for anomaly detection in multivariate time series. Their distinction between abnormal time points, intervals, and series is useful because tenant performance can exhibit all three forms. A single unusual month may be an isolated anomaly; a sequence of abnormal months may indicate a persistent issue; and a tenant whose performance trajectory differs from comparable tenants may represent a portfolio-level anomaly.

Financial-system anomaly research also demonstrates that an anomaly should not automatically be interpreted as fraud. Reviews note that irregularities may arise from deviations, adulterations, or inconsistencies and identify gaps involving data availability, dataset size, and model-effectiveness indicators. This distinction is crucial for tenant auditing. Sales anomalies can arise from legitimate factors such as refurbishment, temporary closure, promotions, product shortages, seasonal demand, or changes in customer behaviour. An intelligent system should therefore produce an investigation signal rather than a misconduct conclusion.

Artificial intelligence research in marketing and retailing points to organisational capability. AI creates value when connected to suitable business processes rather than implemented as technology in isolation. For tenant auditing, model outputs must connect directly to lease administration, audit workflows, and management reporting.

Real-estate research provides the fourth domain. Recent systematic work on AI in real estate indicates that research has expanded substantially, but issues remain around interpretability, data availability, accountability, and organisational readiness. Applications have historically concentrated on valuation, pricing, forecasting, and market analysis, with property management becoming increasingly important. Digital transformation literature similarly identifies operational efficiency and data-driven decision-making as benefits while noting implementation costs, security concerns, resistance to change, and integration difficulties.

The literature therefore reveals a clear gap. Retail analytics provides mature methods for forecasting and pattern recognition, while financial anomaly detection provides techniques for identifying unusual behaviour. Real-estate digitalisation provides the organisational context, but the specific application of these capabilities to tenant

sales auditing under turnover-based leases remains insufficiently developed. There is a need for an integrated framework combining expected-sales modelling, multivariate anomaly detection, peer benchmarking, and lease-related information.

This study responds by positioning tenant sales auditing as a hybrid analytical problem rather than a purely accounting or lease-administration task. The proposed framework treats tenant sales as a time-dependent signal, incorporates contextual variables, identifies deviations from expected performance, and translates those deviations into an explainable audit priority. It also recognises that AI adoption in real estate requires governance, data quality, and human oversight rather than algorithmic automation alone.

3. RESEARCH METHODOLOGY

The study adopts a design-oriented research approach to develop and demonstrate an intelligent tenant sales auditing framework. The approach is appropriate because the objective is to construct a practical analytical architecture that can later be tested with real tenant data. The study combines literature synthesis, conceptual model development, simulated data analysis, and comparative evaluation.

The process contains five stages: business-problem definition; data architecture and simulated dataset preparation; forecasting, anomaly detection, and benchmarking; risk-score generation; and managerial evaluation. The simulated dataset represents 1,200 tenants across eight hypothetical malls and twelve months, producing 14,400 tenant-month records. It includes monthly reported sales, predicted sales, footfall, transaction counts, lease turnover percentage, payment regularity, and peer-group performance. No actual tenant, mall, or proprietary Saudi commercial data are used.

Preprocessing includes missing-value handling, duplicate detection, tenant identifier

matching, scaling of numerical variables, and creation of temporal features. Features include month-on-month sales change, rolling three-month sales average, deviation from category peers, sales-per-visitor ratio, reporting delay, and variance between predicted and reported sales.

Three analytical components are used. First, a forecasting model estimates expected sales from historical tenant performance and contextual variables. Second, an anomaly detector identifies observations that deviate from normal patterns. Third, peer benchmarking compares each tenant with similar tenants within the same category and mall environment. The outputs are combined rather than relying on a single algorithm.

A conceptual anomaly score is calculated as: $\text{Anomaly Score} = |\text{Reported Sales} - \text{Expected Sales}| / \text{Expected Sales}$. A peer deviation score is: $\text{Peer Deviation} = |\text{Tenant Sales} - \text{Median Peer Sales}| / \text{Median Peer Sales}$. A composite risk score is: $\text{Risk Score} = 0.40(\text{Anomaly Score}) + 0.25(\text{Peer Deviation}) + 0.20(\text{Reporting Risk}) + 0.15(\text{Lease/Payment Risk})$. The weights are illustrative and should be calibrated empirically in future research. Scores below 0.33 are classified as low risk, 0.33–0.66 as medium risk, and above 0.66 as high risk.

Model evaluation uses precision, recall, F1-score, mean absolute error for forecasting, and detection time. Because simulated data cannot establish real-world fraud or leakage, the evaluation is a methodological demonstration. The study avoids claiming that the model has proven effectiveness in operating Saudi malls.

Explainability is incorporated. Each alert should show the principal reasons for classification, such as sales variance, abnormal peer deviation, reporting delay, or lease inconsistency. This is important because property managers need defensible explanations before initiating a tenant audit.




Figure 1: Intelligent Tenant Sales Auditing and Anomaly Detection Workflow

Note: Proposed conceptual architecture developed for this study.

Table 1: Core Variables for Intelligent Tenant Sales Auditing

Variable	Type	Example	Purpose
Reported tenant sales	Numerical	Monthly SAR sales	Primary audit input
Predicted sales	Numerical	Model forecast	Expected-performance baseline
Footfall	Numerical	Monthly visitors	Demand context
Peer deviation	Derived	Difference from category median	Benchmarking
Reporting delay	Numerical	Days after deadline	Compliance indicator
Lease turnover rate	Numerical	Contract percentage	Rent context
Payment regularity	Derived	On-time/late	Financial control
Sales trend	Derived	Three-month change	Temporal behaviour

4. Proposed Intelligent Tenant Sales Auditing Framework

4.1 Data Governance and Auditability

The framework is built on the principle that tenant auditing should begin with an expected-performance baseline and investigate deviations that remain unusual after contextual factors are considered. The architecture contains five layers: data acquisition, data preparation, analytical modelling, risk interpretation, and audit action.

The data acquisition layer connects tenant sales declarations, POS feeds where contractually available, lease-management systems, payment records, footfall systems, and portfolio reference data. Not every tenant will provide the same level of information, so the framework is designed for different data-completeness levels. A tenant with direct POS integration may be assessed using transaction-level data, while another may be assessed using monthly declarations and operational indicators.

The data preparation layer creates a common tenant-month record. Data-quality rules detect missing months, duplicate declarations, impossible values, abrupt unit changes, and mismatches between tenant identifiers and lease records.

The analytical layer contains three complementary models. The first is a sales forecasting model. Its role is to estimate a reasonable sales range, not to determine rent directly. Forecast uncertainty should be retained because retail sales are affected by seasonality and unexpected events. A tenant whose reported sales fall slightly below the point forecast but remain within the prediction interval should not necessarily receive a high-risk alert.

The second model is multivariate anomaly detection. It examines relationships between sales, transactions, footfall, historical trends, and peer performance. A sales decline may be normal if footfall

also declines, while a sales decline accompanied by stable footfall and stable peer performance may warrant investigation.

The third component is tenant benchmarking. Each tenant is compared with a relevant peer group defined by mall, retail category, store size where available, and historical operating profile. Benchmarking reduces the risk of comparing fundamentally different businesses.

The risk interpretation layer converts model outputs into a prioritised audit queue. Instead of an opaque alert, the system should identify contributing factors. For example: “High risk: reported sales 27% below model range; sales-per-visitor ratio 31% below peer median; reporting submitted 12 days late.” This allows an auditor to understand the reason for investigation.

The audit action layer completes the feedback loop. Auditors review high-risk cases, determine whether the anomaly is legitimate, and record the outcome. Confirmed discrepancies, legitimate exceptions, and false positives can then be used to recalibrate thresholds and improve future model performance. This creates a human-in-the-loop learning system.

Lease intelligence can provide turnover percentage, reporting deadline, minimum guarantee, and exclusions. Natural language processing may eventually automate extraction from lease documents, but initial implementation can use structured lease-management fields to reduce complexity.

A key design principle is separation between detection and enforcement. An anomaly is a signal, not a finding of wrongdoing. The framework should support audit investigation rather than automatically adjust rent, issue penalties, or accuse a tenant of misreporting. This distinction is essential for procedural fairness, commercial relationships, and responsible AI governance.

Table 2: Proposed AI Components and Audit Functions

Component	Analytical Method	Main Output	Audit Use
Sales forecasting	Time-series/ML forecasting	Expected sales range	Detect under/over-performance
Anomaly detection	Unsupervised/temporal models	Anomaly score	Prioritise unusual cases
Peer benchmarking	Category and mall comparison	Peer deviation	Contextual validation
Lease intelligence	Structured/NLP extraction	Contract indicators	Check reporting/rent rules
Risk scoring	Weighted composite model	Risk category	Create audit queue

This staged approach reduces implementation risk and allows the organisation to demonstrate value before expanding the system. It also creates a governance structure in which technology, leasing, finance, internal audit, information security, and property operations share responsibility for the quality and use of the analytical system.

The roadmap should include model-risk management. Thresholds need periodic review because tenant behaviour, mall composition, consumer preferences, and reporting systems change over time. Model drift can occur when patterns observed during development no longer represent current operations. Monitoring should therefore track changes in prediction error, anomaly rates, tenant mix, and investigation outcomes. Retraining should be triggered by predefined performance or drift thresholds rather than performed arbitrarily.

A practical implementation can proceed through four phases. Phase one establishes data governance, identifies priority malls, maps source systems, and creates the tenant data dictionary. Phase two introduces descriptive dashboards and simple rules so that data-quality problems can be resolved before sophisticated AI models are deployed. Phase three introduces forecasting, anomaly detection, and risk scoring in parallel with existing audits. Phase four scales the framework across the portfolio and introduces continuous model monitoring.

4.2 Model Selection and Validation Strategy

Auditability also requires preservation of the evidence used to create an alert. When the system produces a high-risk classification, the auditor should be able to view the underlying sales series, forecast range, peer comparison, relevant reporting dates, and lease indicators. An immutable audit log can record when the alert was generated, who reviewed it, what evidence was considered, and what conclusion was reached. Such a design improves accountability and creates a valuable training dataset for future model improvement.

Each tenant should receive a persistent identifier that connects the lease record, store location, reporting account, payment account, and

analytical history. Changes in tenant ownership, store relocation, temporary closure, or lease renewal should be recorded as events rather than allowing the model to interpret them as unexplained anomalies. This is particularly important for longitudinal modelling because a structural change in the tenant record can otherwise be mistaken for abnormal commercial behavior.

Data governance is a prerequisite for reliable intelligent tenant auditing. A model can only provide useful signals when the underlying tenant records are complete, correctly matched, and temporally consistent. In practice, sales may be reported at different frequencies, through different channels, or under different definitions. One tenant may provide gross sales while another provides net sales after approved exclusions. The framework should therefore maintain a documented data dictionary defining every variable, unit, reporting period, source system, and transformation rule.

4.3 Tenant Audit Workflow and Human Oversight

The practical value of a model should also be assessed through alert burden. A technically accurate model can still be operationally unsuitable if it generates hundreds of low-value alerts each day. Therefore, management should monitor the percentage of tenant records flagged, the percentage of alerts confirmed after investigation, average review time, and the financial significance of confirmed cases. A useful model is one that improves the yield of audit work rather than simply maximizing a mathematical performance metric.

Model validation should use time-aware procedures rather than random train-test splitting. Random splitting can allow future observations to influence the training set and produce overly optimistic estimates. A rolling-origin evaluation is more appropriate: the model is trained on an initial historical period, tested on a subsequent period, and then updated as additional observations become available. Forecast error can be evaluated using MAE or RMSE, while anomaly detection can be assessed using precision, recall, and F1-score where investigation outcomes are available.

No single machine-learning technique is likely to be optimal for every tenant portfolio.

Forecasting models can perform well when sufficient historical data exist, while unsupervised anomaly detection can be more appropriate where confirmed abnormal labels are scarce. The proposed framework therefore uses a model-portfolio concept. A baseline statistical forecast can establish a transparent reference point, while machine-learning models can capture nonlinear relationships and interactions among sales, footfall, seasonality, and tenant characteristics.

4.4 Portfolio-Level Analytics and Strategic Use

This human-in-the-loop structure reduces the danger of treating algorithmic output as an automatic finding. It also supports continuous improvement because every reviewed alert creates additional information about which signals are useful and which are false positives. Over time, the organization can develop a labelled audit dataset that enables supervised model development and more accurate threshold calibration.

The investigation stage can include requests for supporting sales records, comparison with POS information where contractually available, review of approved exclusions, reconciliation with payment records, and examination of relevant lease clauses. If the anomaly is explained by a legitimate event, the auditor records the explanation. If the discrepancy is confirmed, the financial and contractual resolution follows the organization's established governance procedures. The outcome is then fed back to the analytical system.

The proposed framework creates a structured workflow from detection to investigation. At the beginning of each reporting cycle, tenant data are ingested and quality checks are performed. The forecasting and anomaly models then generate tenant-level indicators. These indicators are combined into a risk score, and the system produces an ordered audit queue. Auditors begin with high-risk cases and review the explanation attached to each alert.

4.5 Implementation Roadmap

The strategic value of the framework therefore extends beyond compliance. Aggregated anomaly information can help asset managers understand the quality of tenant reporting, identify locations requiring stronger data integration, and prioritise investments in digital lease-management infrastructure. Nevertheless, these strategic uses should remain distinct from individual compliance decisions, which require appropriate evidence and human review.

Portfolio analytics can also distinguish isolated tenant events from systemic problems. If

many tenants experience a similar reporting deviation immediately after a system upgrade, the appropriate response may be a technology investigation rather than individual tenant audits. Similarly, if reported sales fall across multiple tenants while footfall remains stable, management can examine whether the change is associated with data integration, reporting definitions, or a broader commercial issue. This capability makes intelligent auditing useful not only for detecting potential discrepancies but also for diagnosing operational weaknesses.

The framework can operate at tenant, mall, category, and portfolio levels. At tenant level, it identifies unusual performance. At mall level, it can detect clusters of anomalies that may indicate a shared operational issue. At category level, it can reveal whether fashion, food and beverage, entertainment, or other segments exhibit systematically different reporting patterns. At portfolio level, management can examine trends in risk exposure and audit outcomes.

5. Simulated Results and Discussion

The simulated evaluation demonstrates how the framework can convert a large tenant portfolio into an actionable audit-priority structure. The simulation assumes 1,200 tenants across eight hypothetical malls and twelve months, producing 14,400 tenant-month observations. The dataset was designed to contain normal seasonal variation as well as synthetic anomalous patterns. These patterns are analytical test cases rather than confirmed real-world leakage events.

The risk model classified approximately 35% of tenants as low risk, 40% as medium risk, and 25% as high risk in the illustrative portfolio. The distribution demonstrates prioritisation: a smaller group can receive intensive audit attention while the majority remains under automated monitoring. The percentages are not estimates of tenant misreporting prevalence in Saudi Arabia.

High-risk cases were characterised by multiple simultaneous deviations. For example, a tenant could show reported sales 24% below its forecast, a sales-per-visitor ratio 28% below its peer group, and repeated reporting delays. The combined score was substantially higher than that generated by any single indicator. This illustrates why a multivariate framework is preferable to a simple rule such as "sales decline above 20% equals high risk."

The forecasting component also demonstrates the importance of uncertainty. Retail sales can fluctuate because of Ramadan, school holidays, promotions, refurbishment, competitor

openings, weather, and macroeconomic conditions. If a model produces only a point forecast, legitimate variation may be incorrectly flagged. A prediction interval or probabilistic forecast can provide a more appropriate basis for auditing.

The illustrative comparison assumes that a conventional manual process detects issues after periodic review, a rule-based process identifies predefined conditions, and the AI-assisted process continuously analyses tenant-month observations. The figures are scenario values used to demonstrate the framework and must not be reported as empirical validation.

The model's value is better understood through its workflow than through a single accuracy percentage. An effective system reduces the number of records requiring human attention while increasing the proportion of reviewed records that contain meaningful anomalies. This is a core audit-triage objective.

The framework also supports portfolio-level pattern detection. If several tenants in one mall simultaneously show lower reported sales while footfall remains stable, a tenant-by-tenant system may generate separate alerts, whereas portfolio analytics can identify a common pattern. Management can then investigate whether the cause is a mall-wide data-integration problem, reporting-system change, payment-system issue, or genuine commercial phenomenon.

Explainability is another important result. The system should provide a ranked list of factors contributing to risk, including revenue variance, peer deviation, reporting delay, and lease indicators. Explainable outputs are important because commercial decisions can affect landlord-tenant relationships and financial settlements.

False positives remain a central concern. A model producing too many alerts may overwhelm audit teams and reduce confidence. Thresholds should therefore be calibrated using historical investigation outcomes. Audit teams should record whether each alert was confirmed, explained, or rejected. These labels can support supervised learning and threshold optimisation.

The framework also creates an opportunity for continuous control monitoring. Instead of treating the annual audit as the main control, management can use monthly or daily analytics to identify emerging risks. This does not eliminate periodic audits; it makes them more targeted.

Overall, the simulated results support the feasibility of an intelligent tenant auditing architecture, while demonstrating why real-world validation is necessary. The principal contribution is a structured method for combining prediction, anomaly detection, benchmarking, and human review in turnover-based lease management.




Figure 2: Illustrative Anomaly-Risk Distribution
Note: Simulated scenario values only; not Saudi market statistics.




Figure 3: Illustrative Comparison of Auditing Approaches

Note: Illustrative scenario assumptions, not empirical results from operating malls.

Table 3: Illustrative Model Evaluation Measures

Measure	Purpose	Interpretation	Use
MAE	Evaluate sales forecasts	Lower is better	Forecast selection
Precision	Evaluate useful alerts	Higher means fewer false positives	Workload control
Recall	Evaluate anomaly coverage	Higher means more anomalies detected	Risk coverage
F1-score	Balance precision and recall	Higher is better	Model comparison
Detection time	Measure operational speed	Lower is better	Process improvement

6. Practical Implications for Saudi Retail Real Estate

The framework has several implications for Saudi retail property operators. First, it can improve audit resource allocation. Large portfolios contain tenants with different commercial profiles, and equal audit intensity is unlikely to be efficient. Risk-based prioritisation allows teams to focus on tenants showing multiple independent indicators of unusual behaviour.

Second, the framework can improve lease-management transparency. When consistent analytical definitions are applied across malls, management can compare reporting quality, anomaly frequency, and audit outcomes. Portfolio dashboards can show risk by mall, category, lease type, or reporting cycle.

Third, the framework can strengthen digital transformation. Saudi retail development increasingly involves digital transaction channels and data-rich customer environments. Integrating these data with lease-management systems creates an opportunity to connect operational intelligence with financial governance. This aligns with broader digital transformation objectives while keeping the research focused on a specific property-management problem.

Fourth, the framework can support commercial decision-making beyond audit. Persistent sales anomalies may reveal underperforming locations, changing consumer preferences, or tenants whose performance differs substantially from peers. Asset managers can use such insights when considering tenant mix, renewal discussions, promotional support, or space reconfiguration. These secondary uses should remain analytically separate from compliance conclusions.

Implementation should occur in stages. A pilot should begin with a limited group of malls and tenants with reliable historical sales data. During the pilot, AI alerts should run in parallel with existing audit procedures. Management can compare alerts with auditor findings and adjust thresholds.

Governance is equally important. Tenant sales data are commercially sensitive, so access controls, audit trails, retention rules, and cybersecurity controls should be established. Model decisions should be explainable, and tenants should not be automatically penalised solely because an algorithm assigns a high-risk score. Human review should remain mandatory for material actions.

The framework can integrate with enterprise systems through APIs connecting POS or

tenant-reporting systems to a central analytics platform. Lease-management software can provide contract metadata, while financial systems can provide payment information. A modular architecture reduces the need to replace existing systems and makes implementation more practical.

7. Limitations and Future Research

The main limitation is the use of simulated data. Simulation allows transparent testing of analytical logic but cannot reproduce all characteristics of actual Saudi tenant sales data. Future research should validate the framework using anonymised historical records from operating malls, including confirmed audit cases and legitimate exceptions.

A second limitation concerns the proposed risk-score weights and thresholds. The weights are conceptual and should not be treated as universally optimal. Future studies can use historical audit outcomes to estimate weights through interpretable statistical and machine-learning methods.

Third, tenant sales are influenced by external variables that may not be captured in basic property datasets. Future research could incorporate promotions, public holidays, weather, macroeconomic indicators, online sales, store opening hours, and mall events. This would improve the ability to distinguish legitimate demand changes from unusual reporting patterns.

Fourth, future studies should compare Isolation Forest, One-Class SVM, autoencoders, and temporal deep-learning models. Evaluation should include false-positive burden, interpretability, computational requirements, and stability over time.

Finally, future research should investigate explainable AI and human-in-the-loop learning. Audit teams should record why alerts are accepted or rejected, allowing the system to learn from professional judgement. Longitudinal studies could assess whether intelligent auditing reduces audit cycle time, improves detection yield, and strengthens lease governance.

8. CONCLUSION

This study developed an intelligent tenant sales auditing framework for turnover-based lease management in Saudi retail real estate. The framework integrates sales forecasting, multivariate anomaly detection, peer benchmarking, lease information, and risk-based audit prioritisation. Its central contribution is the integration of analytical capabilities often treated separately in retail, finance, and property-management systems.

The approach moves tenant auditing from periodic review toward continuous monitoring. It does not assume that every unusual sales pattern represents misreporting; instead, it identifies cases requiring professional investigation. This distinction is essential because legitimate retail volatility can produce anomalies even when reporting is accurate.

The simulated evaluation demonstrates how a large tenant portfolio can be transformed into a structured audit-priority system. However, the results are methodological demonstrations rather than evidence of actual Saudi tenant behaviour. Real-world validation is necessary before the framework can be used for operational or financial decisions.

For Saudi retail property operators, the framework offers a potential pathway toward more transparent, scalable, and data-driven lease management. By combining artificial intelligence with human audit judgement, operators can improve the timing and targeting of tenant reviews while developing stronger portfolio-level visibility. Future empirical research using anonymised commercial data will be necessary to validate model performance, calibrate risk thresholds, and determine the financial and operational benefits of intelligent tenant sales auditing.

REFERENCES

- Chen, Y., Li, X., & Wang, H. (2021). Artificial intelligence-based anomaly detection of energy consumption in buildings: A review, current trends and new perspectives. *Applied Energy*, 287, 116601.
- Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, 48, 24–42.
- Fildes, R., Kolassa, S., & Ma, S. (2022). Post-script—Retail forecasting: Research and practice. *International Journal of Forecasting*, 38(4), 1319–1324. <https://doi.org/10.1016/j.ijforecast.2021.09.012>
- Grewal, D., Hulland, J., Kopalle, P. K., & Karahanna, E. (2021). The future of technology and marketing: A multidisciplinary perspective. *Journal of the Academy of Marketing Science*, 48, 1–8.
- Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications*, 193, 116429. <https://doi.org/10.1016/j.eswa.2021.116429>
- Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science*, 49, 30–50.

- Huang, M.-H., & Rust, R. T. (2022). Artificial intelligence in service. *Journal of Service Research*, 25(1), 3–17.
- Li, G., & Jung, J. J. (2023). Deep learning for anomaly detection in multivariate time series: Approaches, applications, and challenges. *Information Fusion*, 91, 93–113. <https://doi.org/10.1016/j.inffus.2022.10.008>
- Ma, S., & Fildes, R. (2021). Retail sales forecasting with meta-learning. *European Journal of Operational Research*, 288(1), 111–128. <https://doi.org/10.1016/j.ejor.2020.05.038>
- Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). M5 accuracy competition: Results, findings, and conclusions. *International Journal of Forecasting*, 38(4), 1346–1364.
- Policarpo, L. M., da Silveira, D. E., Righi, R. R., Stoffel, R. A., da Costa, C. A., Barbosa, J. L. V., Scorsatto, R., & Arcot, T. (2021). Machine learning through the lens of e-commerce initiatives: An up-to-date systematic literature review. *Computer Science Review*, 41, 100414. <https://doi.org/10.1016/j.cosrev.2021.100414>
- Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and artificial intelligence: An experiential perspective. *Journal of Marketing*, 85(1), 131–151.
- Ruff, L., Kauffmann, J. R., Vandermeulen, R. A., Montavon, G., Samek, W., Kloft, M., Dietterich, T. G., & Müller, K.-R. (2021). A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, 109(5), 756–795.
- Thudumu, S., Branch, P., Jin, J., & Singh, J. J. (2020). A comprehensive survey of anomaly detection techniques for high dimensional big data. *Journal of Big Data*, 7, 61.
- Villa-Pérez, M. E., Álvarez-Carmona, M. Á., Loyola-González, O., Medina-Pérez, M. A., Velazco-Rossell, J. C., & Choo, K.-K. R. (2021). Semi-supervised anomaly detection algorithms: A comparative summary and future research directions. *Knowledge-Based Systems*, 218, 106878.
- Verma, S., Sharma, R., Deb, S., & Maitra, D. (2021). Artificial intelligence in marketing: Systematic review and future research direction. *International Journal of Information Management Data Insights*, 1(1), 100002. <https://doi.org/10.1016/j.ijime.2020.100002>
- Wamba-Taguimdje, S.-L., Fosso Wamba, S., Kamdjoug, J. R. K., & Tchatchouang Wanko, C. E. (2020). Influence of artificial intelligence on firm performance: The business value of AI-based transformation projects. *Business Process Management Journal*, 26(7), 1893–1924.
- Da, Z., Dun, Y., Liu, C., Liang, Y., Xue, Y., & Qian, X. (2023). Anomaly detection framework for unmanned vending machines. *Knowledge-Based Systems*, 262, 110251.
- Patil, S., Varadarajan, V., Mazhar, S. M., Sahibzada, A., Ahmed, N., Sinha, O., Kumar, S., Shaw, K., & Kotecha, K. (2022). Explainable artificial intelligence for intrusion detection system. *Electronics*, 11(19), 3079.
- Fosić, I., Žagar, D., Grgić, K., & Križanović, V. (2023). Anomaly detection in NetFlow network traffic using supervised machine learning algorithms. *Journal of Industrial Information Integration*, 33, 100466.
- Mariani, M. M., Perez-Vega, R., & Wirtz, J. (2022). AI in marketing, consumer research and psychology: A systematic literature review and research agenda. *Psychology & Marketing*, 39(4), 755–776.
- Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., et al. (2021). Artificial intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57, 101994.
- Shankar, V. (2021). How artificial intelligence is reshaping retailing. *Journal of Retailing*, 97(1), 6–12.
- Kaplan, A., & Haenlein, M. (2020). Rulers of the world, unite! The challenges and opportunities of artificial intelligence. *Business Horizons*, 63(1), 37–50.
- Knight Frank. (2025, June 1). Riyadh leads Saudi retail rebound as consumer spending surges to SAR 1.41 trillion. Knight Frank Saudi Arabia. <https://www.knightfrank.com.sa/en/newsroom/article/2025/6/ksa-retail-market-review-spring-2025>
- Knight Frank. (2025, November 10). The Saudi Report 2025: Retail & F&B. Knight Frank Saudi Arabia.
- National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0). U.S. Department of Commerce.
- OECD. (2024). OECD artificial intelligence principles and policy considerations. OECD Publishing.
- World Economic Forum. (2020). The future of jobs report 2020. World Economic Forum.