

AI-Enabled Population Health Analytics for Predicting Chronic Disease Risk and Preventable Hospital Admissions in Saudi Arabia

Salah Ali Ibrahim Ramadan^{1*}

Independent Researcher

*Corresponding Author

Salah Ali Ibrahim Ramadan

Independent Researcher

Article History

Received: 19.08.2026

Accepted: 12.10.2026

Published: 14.10.2026

Abstract: Saudi Arabia's shift towards prevention, integrated care and digital health creates an opportunity to use artificial intelligence at population scale rather than only for isolated diagnosis. This review examines how AI-enabled population health analytics can combine clinical, administrative, insurance, pharmacy, laboratory and patient-generated data to predict chronic disease risk and preventable hospital admissions. Thirty peer-reviewed studies published from 2020 to 2025 were synthesised across Saudi health-system transformation, digital maturity, chronic disease prediction, electronic records, remote monitoring, risk stratification, readmission prediction, explainable AI and avoidable hospitalisation. The evidence shows that machine-learning models can identify elevated risk for diabetes, cardiovascular disease, sepsis, emergency use and readmission, but predictive performance alone does not create prevention. Population benefit depends on reliable longitudinal data, clinically meaningful prediction horizons, transparent risk thresholds, human review and direct links to primary-care outreach, screening, medication review, case management and remote monitoring. Claims and encounter data provide breadth, while laboratory values, vital signs, clinical notes and patient-reported information improve clinical specificity. The review proposes a six-layer Population Health Intelligence Framework comprising multi-source data integration, longitudinal patient resolution, AI risk modelling, clinical validation, preventive care orchestration and outcome-based model learning. Preventable admissions should be treated as a system outcome requiring both accurate risk detection and timely access to effective primary care. For Saudi Arabia, the framework aligns predictive analytics with Vision 2030 goals for prevention, efficiency and patient-centred care while emphasising interoperability, privacy, fairness, explainability and accountable clinical governance.

Keywords: Population Health Analytics, Artificial Intelligence, Chronic Disease, Preventable Hospital Admissions, Risk Stratification, Saudi Arabia, Primary Care, Machine Learning, Vision 2030.

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 faces a growing burden of chronic disease at the same time as its health system is being reorganised around prevention, digital transformation and more integrated models of care.

Primary healthcare reform has expanded preventive functions and care coordination, while Vision 2030 has accelerated investment in digital records, virtual care, analytics and population health management [1-29]. This creates a policy environment in which

Citation: Ramadan, S. A. I. (2026). AI-Enabled Population Health Analytics for Predicting Chronic Disease Risk and Preventable Hospital Admissions in Saudi Arabia; *Glob Acad J Econ Buss*, 8(5), 1840-1850. <https://doi.org/10.36348/gajeb.2026.v08i05.069>

artificial intelligence can move beyond narrow diagnostic tasks and support proactive identification of people whose risk is increasing before complications require hospital treatment.

The distinction between individual prediction and population health intelligence is important. A disease classifier may identify whether one patient has a high probability of diabetes, cardiovascular disease or another condition. Population health analytics must go further: it should identify cohorts, detect changing risk, prioritise limited clinical resources, reveal gaps in preventive care, and monitor whether interventions reduce emergency use and avoidable admission. This requires integration across clinical and administrative data rather than dependence on a single model or source.

Saudi research demonstrates the technical feasibility of predictive analytics across several domains. Machine-learning methods have been used to classify cardiovascular risk, diabetes prevalence, high-risk type 2 diabetes, sepsis and emergency-department data [5-10]. Remote monitoring and connected-health architectures can add continuous physiological information to episodic records [11-13]. These studies establish a useful evidence base, but they also expose recurring limitations: single-centre datasets, historical data, class imbalance, incomplete variables, weak external validation and limited connection between model output and preventive workflows.

Preventable hospital admissions provide a more demanding test of value. Admission for uncontrolled hypertension, diabetes, heart failure, chronic respiratory disease or other ambulatory-care-sensitive conditions may reflect disease severity, but it can also indicate delayed detection, weak continuity, medication problems, limited access or failure to act on worsening risk. Recent Saudi evidence on hypertension in Makkah illustrates the value of tracking preventable hospitalisation alongside primary-care use as a health-system performance signal [25]. Predictive analytics can potentially identify rising risk earlier, yet the model must be coupled with an intervention capable of changing the outcome.

Risk stratification literature reinforces this point. Claims, utilisation and social data can improve identification of high-cost or high-need populations, but claims-only models can become stale and may miss social or clinical context that professionals know about [21, 22]. Electronic records improve granularity, while patient-generated and remote-monitoring data can add behavioural and physiological signals. The challenge is to create a

trusted longitudinal patient view in which different sources are reconciled, time-aligned and interpreted according to their original purpose.

The second challenge is actionability. A model that predicts a 30-day admission risk of 25% is useful only if clinicians know why risk is elevated, whether the signal is reliable and what can still be changed. For one patient, the appropriate response may be medication review; for another, an urgent primary-care appointment, home monitoring, specialist assessment or social support. Population health systems must therefore translate risk scores into differentiated care pathways rather than generic alerts.

The third challenge is governance. Healthcare AI affects access, prioritisation and clinical attention. Models may perform differently across age groups, regions, sexes or disease subgroups, and poor data quality can create systematic under-identification [18]. Explainability, transparency and human review are therefore not optional additions but core design requirements [19, 20]. Saudi implementation must also address record interoperability, privacy, workforce readiness and digital maturity [3-15].

This review develops an integrated framework for AI-enabled population health analytics focused specifically on chronic disease risk and preventable hospital admissions. It synthesises Saudi and international evidence to connect data architecture, prediction, care orchestration and outcome measurement into one preventive operating model.

2. Aim, Objectives and Review Questions

The aim of this review is to develop an evidence-informed framework for using AI-enabled population health analytics to predict chronic disease risk and preventable hospital admissions in Saudi Arabia.

Four objectives guide the review. First, it examines which health data sources provide useful signals for population-level chronic disease prediction. Second, it evaluates how machine-learning and related predictive methods can support risk stratification while preserving clinical interpretability. Third, it identifies preventive care pathways that can convert predicted risk into action. Fourth, it defines governance, measurement and implementation requirements for using predictive analytics safely across Saudi healthcare.

Three review questions structure the synthesis. RQ1 asks how clinical, administrative, insurance, pharmacy, remote-monitoring and

patient-generated data should be combined for longitudinal risk prediction. RQ2 asks which analytical and governance features make a risk model clinically actionable rather than merely accurate. RQ3 asks how predictive risk should be connected to primary care and care-management interventions to reduce preventable hospital use.

3. REVIEW METHODOLOGY

A structured integrative review was selected because the topic spans population health, clinical informatics, machine learning, chronic disease management, primary care, health-system operations and AI governance. The evidence window was restricted to 2020-2025. Peer-reviewed studies were prioritised where they contributed directly to Saudi healthcare transformation, digital readiness, chronic disease prediction, risk stratification, electronic records, remote monitoring, admission or readmission prediction, or implementation governance.

Search concepts combined artificial intelligence, machine learning, population health, chronic disease, diabetes, cardiovascular disease, hypertension, hospital admission, readmission, avoidable hospitalisation, primary care, electronic health records, claims, risk stratification and Saudi Arabia. International studies were retained when they described mechanisms that could reasonably transfer to the Saudi context, risk-scoring implementation, electronic-record value, readmission prediction and avoidable hospitalisation.

The final evidence set contained 30 references, all published within the requested period and each linked to a resolvable DOI. Evidence was coded under six domains: data foundation, risk modelling, chronic disease detection, hospital-use prediction, preventive intervention and governance. A second interpretive pass identified recurring mechanisms including record linkage, temporal alignment, feature engineering, calibration, threshold design, clinician override, social-risk supplementation, care-gap detection, case management and continuous recalibration.

Because the included studies differ substantially in populations, algorithms, endpoints and reported metrics, numerical results were not pooled. The review instead compares the operational logic behind successful predictive systems. No primary data were created and no causal effect sizes were invented. Technical performance is discussed only in relation to whether models could support real-world preventive decisions.

4. Population Health Data Foundation

Population health prediction depends first on the quality of the longitudinal record. Saudi hospitals and primary-care organisations use electronic records, but digital maturity and interoperability remain uneven [3,4]. A population model should therefore be designed around data provenance and linkage rather than assuming that all available records describe the same patient state with equal reliability.

Clinical records contribute diagnoses, vital signs, laboratory results, medication lists and narrative assessments [17-23]. Administrative encounter data contribute utilisation history, admissions, emergency attendance, procedures and service locations. Insurance and claims records provide a broad longitudinal view of utilisation, eligibility and expenditure. Pharmacy data reveal dispensing patterns and potential adherence gaps. Remote-monitoring systems contribute frequent physiological signals, while patient-reported information can add symptoms, lifestyle and social context not routinely represented in billing records.

Each source has a different temporal meaning. A claim confirms that a billed service occurred but may appear after a delay. A laboratory result can indicate current physiological control. A diagnostic code may represent established disease, screening, rule-out coding or historical history. A wearable device can provide near-real-time measurements but may be affected by adherence or device quality. Models should therefore preserve timestamps, source identity and uncertainty rather than collapsing all observations into an undifferentiated feature set.

Identity resolution is important. Longitudinal prediction fails when one person's data are fragmented across systems or when duplicate records are mistakenly treated as separate patients. A trusted patient layer should link records through governed identifiers, reconcile terminology and maintain traceability back to the source. Standardised disease, laboratory, medication and procedure vocabularies reduce variation that otherwise appears to the algorithm as clinical difference.

Missingness should be interpreted carefully. The absence of a laboratory value may indicate that a test was unnecessary, unavailable, performed elsewhere or simply not captured. Treating missingness as zero can create misleading patterns. Similarly, low utilisation does not necessarily mean low need; people facing geographic, financial or behavioural access barriers may appear healthier in utilisation data than they actually are. Risk

stratification research shows that social context and professional knowledge can correct misclassification produced by administrative data alone [21, 22].

Table 1 summarises the principal data domains, their contribution to prediction and their limitations.

Table 1: Population-health data domains and analytical value

Data domain	Illustrative content	Predictive contribution	Key limitation
Primary-care EHR	Diagnoses, vitals, laboratory results, clinical notes	Current disease state and control	Fragmentation; missing values; coding variation
Hospital encounters	Admissions, ED visits, procedures, discharge information	Acute utilisation and prior deterioration	Hospital-centred view; may be late for prevention
Claims/insurance	Longitudinal utilisation, eligibility, procedures, costs	Population segmentation and utilisation history	Coding lag; limited clinical granularity
Pharmacy	Prescriptions, dispensing, refill continuity	Medication gaps and treatment intensity	Dispensing does not prove adherence
Remote monitoring	Blood pressure, glucose, activity, wearable signals	Trajectory and early physiological change	Device adherence; noise; digital access
Patient/social data	Symptoms, lifestyle, function, social determinants	Context, self-management and barriers to care	Response bias; uneven capture; privacy sensitivity

5. Chronic Disease Risk Prediction

Chronic disease prediction is one of the strongest areas of Saudi AI evidence. Cardiovascular disease research has used demographic, behavioural and physiological inputs to classify risk, showing that machine-learning and neuro-fuzzy methods can support earlier identification [5]. Diabetes research has similarly used behavioural factors, coded clinical records and routinely collected variables to estimate prevalence or identify individuals at elevated risk [6-30].

These studies reveal several design lessons. First, the prediction target must be clinically precise. "Diabetes risk" may refer to prevalent undiagnosed disease, future onset, poor glycaemic control or complications. Each target requires different labels and intervention timing. A population health platform should therefore define whether the model predicts onset, deterioration or acute utilisation before choosing features or thresholds.

Second, predictive performance must be evaluated beyond overall accuracy. Chronic disease datasets are often imbalanced; a model can appear accurate while missing a substantial proportion of high-risk patients. Sensitivity, specificity, positive predictive value, calibration and subgroup performance are important because the consequences of false negatives and false positives differ. Screening programmes may prefer higher sensitivity, whereas intensive case-management programmes may require higher precision because capacity is limited.

Third, external validation matters. Several Saudi studies use data from one institution, historical period or specialised population [5-10]. A model trained in a tertiary centre may not generalise to

primary care, another region or a different demographic mix. Population deployment should therefore include geographical and temporal validation before a score is used for care prioritisation.

Fourth, models should focus on modifiable risk where possible. Age and prior disease are powerful predictors but cannot be changed. From a preventive perspective, the highest-value features are those that reveal a pathway for action: uncontrolled blood pressure, rising HbA1c, medication gaps, repeated emergency use, declining activity or missed follow-up. The objective is not only to identify who is likely to deteriorate, but to identify who may benefit from intervention.

Remote monitoring expands this capability. Connected-health and Internet-of-Things architectures can stream vital signs and behavioural data into analytic systems [11-13]. For hypertension, diabetes and heart failure, repeated measurements can reveal trajectory rather than a single clinic value. However, data volume should not be confused with clinical value. The monitoring programme needs thresholds, confirmation logic and a response pathway to avoid overwhelming clinicians with low-value alerts.

6. Predicting Hospital Use and Preventable Admissions

Hospital-use prediction differs from disease prediction because admission is determined by both clinical severity and health-system factors. A person with poorly controlled hypertension may remain out of hospital if primary care identifies deterioration early, medication is adjusted and follow-up is accessible. Another person with similar physiology may be admitted because care is delayed.

Preventable-admission prediction must therefore include both patient risk and care-process signals.

Emergency-department analytics in Saudi Arabia demonstrates how machine learning can classify and monitor patient flows using clinical history, triage and length-of-stay variables [10]. ICU prediction studies show that algorithms can identify patients at high risk of deterioration or mortality [7, 8]. These acute-care applications are useful, but population prevention requires shifting the prediction horizon earlier, before the patient enters the hospital pathway.

Readmission models provide a closer analogue. Reviews show that random forests, neural networks and other machine-learning methods can improve 30-day readmission risk prediction compared with traditional approaches in some settings [24]. Yet readmission risk often depends on information that is poorly represented in structured data: functional status, caregiver support, medication understanding, access to follow-up and social instability. The strongest population models therefore combine clinical severity with utilisation history and contextual variables.

Avoidable hospitalisation metrics can provide a system-level outcome. International work on ambulatory-care-sensitive conditions shows that certain admission categories can be used to compare primary-care performance when definitions are carefully harmonised [26]. In Saudi Arabia, recent analysis of hypertension-related preventable hospitalisation in Makkah provides direct evidence that avoidable admission can be monitored alongside primary-care activity [25]. This establishes a relevant outcome for evaluating whether population analytics and preventive programmes are working.

Prediction should distinguish several horizons. Near-term admission risk, such as seven or thirty days, supports intensive outreach and transition management. Medium-term risk may support chronic-disease review, medication optimisation and specialist referral. Longer-term disease progression risk supports screening, lifestyle programmes and resource planning. Mixing these horizons into one score can make intervention unclear.

A useful preventable-admission model should also distinguish inevitability from modifiability. Some admissions are clinically appropriate despite excellent outpatient care. The objective is not to suppress necessary hospital use, but to identify admissions that may be avoided through earlier intervention. This requires clinical review of outcome definitions and caution when

interpreting reduced admission counts as automatic evidence of improved care.

7. Risk Stratification and Actionability

Risk stratification translates continuous model outputs into operational categories. A practical population system can use several tiers: routine prevention, emerging risk, high risk and complex/high-acuity management. The purpose of the tier is to match intervention intensity to expected benefit and available resources.

Claims and utilisation data are valuable for broad stratification because they cover large populations and capture prior service use. However, implementation research shows that claims-based scores can be outdated or incomplete and may miss socially situated risks known to clinicians [21, 22]. Clinical data, admission-discharge-transfer signals and social determinants can improve relevance, but additional variables also increase complexity and governance requirements.

Thresholds should be set according to intervention capacity and clinical consequence rather than statistical optimisation alone. If a primary-care programme can contact five thousand people per month, a threshold that identifies fifty thousand will create an unmanaged queue. Conversely, a very restrictive threshold may miss people whose disease is worsening but still highly modifiable. Capacity-aware threshold design is therefore a core population-health function.

Risk explanations should accompany the score. Clinicians need to know whether elevated risk is driven by uncontrolled blood pressure, repeated emergency attendance, declining renal function, medication gaps or another factor. Explanations facilitate validation and help select the right intervention. They also support patient communication when a predictive programme influences follow-up frequency or referral.

Human review is especially important for high-consequence decisions. Algorithms can prioritise records for review, but they should not automatically deny services or determine that hospital care is unnecessary. Explainable and transparent AI frameworks are essential for trust, particularly when models are embedded in routine clinical workflows [19, 20].

8. Preventive Care Orchestration

Prediction creates value only when linked to a care pathway. The proposed framework therefore treats preventive orchestration as a distinct layer between analytics and outcomes. Once a high-risk patient is identified, the system should determine

who owns the case, what action is recommended, how quickly it should occur and how completion will be documented.

For emerging chronic disease risk, interventions may include laboratory screening, blood-pressure review, lifestyle counselling, medication reconciliation or appointment support. For established disease with evidence of deterioration, interventions may include nurse outreach, remote monitoring, pharmacist review or expedited primary-care assessment. For complex multimorbidity, multidisciplinary case management and specialist coordination may be required.

Care-gap detection can operate alongside risk prediction. A patient with diabetes who has missed renal monitoring or retinal screening represents an actionable gap even if the model does not predict imminent admission. Combining risk and care gaps helps distinguish urgency from preventive opportunity. Saudi work on diabetic retinopathy illustrates how AI-supported screening can extend preventive reach when embedded within a clinically governed pathway [14].

Closed-loop documentation is necessary. A population health record should capture the original signal, review decision, intervention, contact attempts, referrals, completion and outcome. This provides accountability and generates learning data. If high-risk alerts repeatedly lead to no action, the problem may be workflow capacity rather than model performance. If interventions are completed but admissions do not fall, the target, timing or intervention may need revision.

Digital and remote care can increase reach. Telehealth and connected-monitoring systems can support follow-up without requiring every interaction to occur in hospital [12-16]. This is particularly relevant for chronic disease, where repeated low-intensity contact may prevent episodic deterioration. However, digital access should be monitored to avoid widening disparities between populations with different connectivity, digital literacy or language needs.

9. Integrated Population Health Intelligence Framework

The proposed framework contains six layers. Layer one collects population health data from primary care, hospitals, claims, pharmacy, laboratory systems, remote monitoring and patient-reported sources. Layer two creates a trusted longitudinal patient record through identity resolution, terminology mapping, temporal alignment, provenance and quality controls. Layer three applies AI models to disease onset, disease-control deterioration, emergency use, admission and readmission risk. Layer four performs clinical validation, applies governed thresholds and reviews explainable drivers. Layer five converts risk into preventive care actions. Layer six measures outcomes and recalibrates the system.

Figure 1 presents this architecture as a continuous pathway from data to prevention. Figure 2 emphasises the population perspective by connecting integrated data sources to modelling, intervention and health-system outcomes. The framework is intentionally closed loop because the effectiveness of a risk model depends on what happens after prediction.




Figure 1: Six-layer population health intelligence architecture from multi-source data to preventive action and continuous learning




Figure 2: Population-level analytics framework linking integrated Saudi health data, risk modelling, intervention and health-system outcomes

Four distinctions are central. First, surveillance is different from prediction: prevalence and burden dashboards support planning, whereas individual risk scores support intervention. Second, prediction is different from causation: a variable associated with admission risk is not necessarily a suitable treatment target. Third, risk is different from preventability: some high-risk admissions remain

clinically necessary. Fourth, model performance is different from programme performance: a well-calibrated model can fail to improve outcomes if the care pathway is weak.

Table 2 translates these principles into governance gates and performance measures.

Table 2: Governance gates and performance measures for population health analytics

Gate	Core requirements	Illustrative measures
Data gate	Identity resolution; provenance; terminology mapping; timestamp alignment	Duplicate rate; missingness; freshness; mapping completeness
Model gate	Defined target and horizon; calibration; subgroup validation; explainability	AUROC plus calibration; sensitivity; precision; subgroup error
Clinical gate	Human confirmation; modifiable drivers; override with rationale	Review time; override yield; false-positive burden
Intervention gate	Risk-tier pathway; named owner; service capacity; closed-loop referral	Contact success; screening completion; referral closure
Outcome gate	Disease control; ED use; preventable admission; readmission; patient outcomes	HbA1c/BP control; avoidable admissions; 30-day readmission
Learning gate	Drift monitoring; outcome feedback; threshold and model recalibration	Drift; calibration change; intervention effectiveness; equity gaps

10. Outcome Measurement and Evaluation

Evaluation should separate model metrics from population outcomes. Model metrics include discrimination, calibration, sensitivity, specificity, precision and error. Workflow metrics include time from signal to review, contact completion, referral closure and intervention uptake. Clinical outcomes include blood-pressure control, glycaemic control, complication rates and patient-reported status. Utilisation outcomes include emergency visits, preventable admissions, readmissions and length of stay.

The evaluation period should match the prediction horizon. Screening programmes may require months or years before disease incidence changes, whereas a post-discharge readmission programme can be assessed within thirty days. Leading indicators such as appointment completion or medication reconciliation should therefore be paired with lagging outcomes.

Fairness should be monitored by subgroup. Performance can differ across region, age, sex, nationality, socioeconomic context or disease complexity. A model with acceptable average calibration may still systematically underpredict risk

in a vulnerable subgroup. Equity dashboards should therefore track selection rates, false negatives, intervention completion and outcomes across relevant populations.

Outcome attribution requires caution. A decline in admissions may occur because of changes in coding, hospital capacity, disease prevalence or another programme. Evaluation designs should be used when organisations claim clinical or economic benefit. Routine dashboards are suitable for operational monitoring, but causal claims may require matched comparisons, interrupted time-series analysis or prospective evaluation.

11. Governance, Privacy and Trust

Population analytics concentrates sensitive information across organisational boundaries. Privacy, access control and auditability must therefore be built into the architecture. Data should be used for defined clinical and population-health purposes, with minimum-necessary access and documented provenance. Security mechanisms for electronic health records and distributed data environments are directly relevant to Saudi implementation [15].

Interoperability is both technical and organisational. A model cannot be trusted if diagnoses, laboratory units or admission statuses are defined differently across facilities. Common data definitions and quality rules should be governed nationally or at network level where population models cross provider boundaries.

Model governance should document intended use, population, outcome definition, prediction horizon, threshold logic, performance, limitations and recalibration schedule. Drift monitoring is essential because disease prevalence, clinical practice, coding and service access can change over time. A model trained before a major care redesign may become miscalibrated afterwards.

Workforce capability is equally important. Saudi digital transformation studies identify predictive analytics, interoperability and workforce readiness as areas requiring continued development [3, 4]. Clinicians need to understand what a risk score can and cannot say; analysts need clinical context; and leaders need mechanisms for acting on early warnings without creating unsafe automation [27].

12. Implementation Roadmap for Saudi Arabia

Implementation should begin with a focused disease pathway and a clearly defined avoidable outcome. Hypertension and diabetes are strong candidates because they are prevalent, measurable and closely connected to primary-care prevention. A

pilot could identify rising-risk patients, assign outreach, track completion and measure emergency use or preventable admission over time.

Stage one establishes data linkage, terminology and quality standards. Stage two creates descriptive cohort surveillance and care-gap measures. Stage three validates transparent baseline risk models. Stage four adds more complex machine learning only when it improves calibration or actionable identification. Stage five embeds risk tiers into primary-care and case-management workflows. Stage six measures outcomes, equity and workload, then recalibrates models and thresholds.

A decision baseline should be recorded before deployment. Existing admission rates, emergency use, screening completion, follow-up delays and workload should be measured so improvement can be demonstrated rather than assumed. Capacity planning should accompany analytics; identifying more high-risk patients without expanding preventive services can worsen waiting times and clinician burden.

A practical implementation should also include a shadow-deployment phase before model outputs influence care prioritisation. During this period, predictions can be generated in parallel with existing clinical workflows and compared with clinician judgement, subsequent utilisation and actual disease-control outcomes. This allows teams to identify unstable features, inappropriate thresholds and unexpected workload consequences without changing patient care prematurely. Alert burden should be measured explicitly, because excessive low-value warnings can reduce attention to genuinely urgent cases. The pilot should also examine reasons for clinician override. Repeated overrides may indicate weak calibration, missing contextual variables or disagreement about the intended use of the model. Conversely, high acceptance does not automatically prove clinical benefit; teams must still verify whether completed interventions improve outcomes. A controlled deployment process therefore links technical validation with operational validation. It also creates an auditable record of how predictive information is interpreted, which is important when risk scores affect appointment priority, case-management intensity or referral. Scaling across regions should occur only after performance, equity, workload and intervention capacity have been demonstrated under real operating conditions and after local teams confirm that the model complements rather than disrupts established clinical accountability. for all patients.

The implementation model is therefore multidisciplinary. Population-health clinicians define

modifiable risk and pathways; data teams build reliable features; primary-care teams deliver interventions; pharmacists support medication optimisation; care managers coordinate complex cases; and governance teams monitor privacy, fairness and safety.

13. Research Priorities

Future Saudi research should move from single-centre model development towards multicentre prospective validation. Models should be tested across regions, public and private settings, and different levels of care. Calibration and subgroup performance should be reported alongside discrimination.

Research should also connect prediction with intervention. It is not enough to show that a model identifies high-risk diabetes or cardiovascular patients. Studies should test whether model-guided outreach improves screening, disease control or preventable admission compared with usual care.

A third priority is richer longitudinal data. Claims and electronic records should be supplemented, where appropriate, with pharmacy, remote-monitoring and patient-reported information. Research should examine which additional variables meaningfully improve decisions rather than simply improving statistical fit.

Finally, economic evaluation is needed. Population analytics requires infrastructure, validation, clinical review and preventive capacity. Studies should examine whether earlier intervention reduces avoidable acute-care costs sufficiently to justify these investments, while also measuring patient experience and equity.

14. CONCLUSION

AI-enabled population health analytics can support Saudi Arabia's transition from reactive treatment towards earlier detection and prevention, but its value depends on an integrated clinical operating model. Saudi studies demonstrate that machine learning can identify chronic disease and acute deterioration risk, while international evidence shows how multi-source risk stratification and readmission prediction can support population management. The central challenge is not generating a score; it is converting reliable risk information into timely preventive action.

The proposed framework links multi-source data, longitudinal patient resolution, AI prediction, clinical validation, preventive care orchestration and continuous outcome learning. It recognises that chronic disease risk, hospital-use risk and preventability are related but distinct concepts.

Models should therefore use explicit prediction horizons, transparent thresholds and explainable drivers, with human review for consequential decisions.

For Saudi Arabia, preventable hospital admissions offer a meaningful outcome for evaluating whether digital intelligence is improving care beyond technical model performance. When predictive analytics is connected to primary care, medication management, remote monitoring and case coordination, the health system gains time to intervene before deterioration becomes an acute event.

The strategic benefit under Vision 2030 is therefore a stronger learning health system: one that identifies risk earlier, targets preventive resources more precisely, measures whether interventions work, detects inequity and continually recalibrates as population needs change. That capability can improve chronic disease control and reduce avoidable hospital use while preserving clinical accountability and patient trust.

REFERENCES

1. Al Khashan H, Abogazalah F, Alomary S, et al. Primary health care reform in Saudi Arabia: progress, challenges and prospects. *Eastern Mediterranean Health Journal*. 2021;27(10):1016-1026. <https://doi.org/10.26719/emhj.21.042>
2. Mani ZA, Goniewicz K. Transforming healthcare in Saudi Arabia: a comprehensive evaluation of Vision 2030's impact. *Sustainability*. 2024;16(8):3277. <https://doi.org/10.3390/su16083277>
3. Al-Kahtani N, Alruwaie S, Al-Zahrani BM, et al. Digital health transformation in Saudi Arabia: a cross-sectional analysis using Healthcare Information and Management Systems Society digital health indicators. *Digital Health*. 2022;8:20552076221117742. <https://doi.org/10.1177/20552076221117742>
4. Alhur A. Overcoming Electronic Medical Records Adoption Challenges in Saudi Arabia. *Cureus*. 2024;16(2):e53827. <https://doi.org/10.7759/cureus.53827>
5. Taylan O, Alkabaa AS, Alqabbaa HS, Pamukcu E, Leiva V. Early Prediction in Classification of Cardiovascular Diseases with Machine Learning, Neuro-Fuzzy and Statistical Methods. *Biology*. 2023;12(1):117. <https://doi.org/10.3390/biology12010117>
6. Almutairi ES, Abbod MF. Machine Learning Methods for Diabetes Prevalence Classification in Saudi Arabia. *Modelling*. 2023;4(1):37-55. <https://doi.org/10.3390/modelling4010004>

7. Alanazi A, Aldakhil L, Aldhoayan M, Aldosari B. Machine Learning for Early Prediction of Sepsis in Intensive Care Unit Patients. *Medicina*. 2023;59(7):1276. <https://doi.org/10.3390/medicina59071276>
8. Elhazmi A, Al-Omari A, Sallam H, et al. Machine learning decision tree algorithm role for predicting mortality in critically ill adult COVID-19 patients admitted to the ICU. *Journal of Infection and Public Health*. 2022;15(7):826-834. <https://doi.org/10.1016/j.jiph.2022.06.008>
9. Alkattan A, Al-Zeer A, Alsaawi F, et al. The utility of a machine learning model in identifying people at high risk of type 2 diabetes mellitus. *Expert Review of Endocrinology & Metabolism*. 2024;19(6):513-522. <https://doi.org/10.1080/17446651.2024.2400706>
10. Ragab M, Kateb F, Al-Rabia MW, et al. A Machine Learning Approach for Monitoring and Classifying Healthcare Data-A Case of Emergency Department of KSA Hospitals. *International Journal of Environmental Research and Public Health*. 2023;20(6):4794. <https://doi.org/10.3390/ijerph20064794>
11. Qaffas AA, Hoque R, Almazmomi N. The Internet of Things and Big Data Analytics for Chronic Disease Monitoring in Saudi Arabia. *Telemedicine and e-Health*. 2021;27(1):74-81. <https://doi.org/10.1089/tmj.2019.0289>
12. Ahila A, Dahan F, Alrooba R, et al. A smart IoMT based architecture for E-healthcare patient monitoring system using artificial intelligence algorithms. *Frontiers in Physiology*. 2023;14:1125952. <https://doi.org/10.3389/fphys.2023.1125952>
13. Alnaim AK, Alwakeel AM. Machine-Learning-Based IoT-Edge Computing Healthcare Solutions. *Electronics*. 2023;12(4):1027. <https://doi.org/10.3390/electronics12041027>
14. Barakat AA, Mobarak O, Javaid HA, et al. The application of artificial intelligence in diabetic retinopathy screening: a Saudi Arabian perspective. *Frontiers in Medicine*. 2023;10:1303300. <https://doi.org/10.3389/fmed.2023.1303300>
15. Alruwaili FF. Artificial intelligence and multi agent based distributed ledger system for better privacy and security of electronic healthcare records. *PeerJ Computer Science*. 2020;6:e323. <https://doi.org/10.7717/peerj-cs.323>
16. El-Sherif DM, Abouzid M, Elzarif MT, et al. Telehealth and Artificial Intelligence Insights into Healthcare during the COVID-19 Pandemic. *Healthcare*. 2022;10(2):385. <https://doi.org/10.3390/healthcare10020385>
17. Sim JA, Huang X, Horan MR, et al. Natural language processing with machine learning methods to analyze unstructured patient-reported outcomes derived from electronic health records: a systematic review. *Artificial Intelligence in Medicine*. 2023;146:102701. <https://doi.org/10.1016/j.artmed.2023.102701>
18. Esmaeilzadeh P. Challenges and strategies for wide-scale artificial intelligence deployment in healthcare practices: a perspective for healthcare organizations. *Artificial Intelligence in Medicine*. 2024;151:102861. <https://doi.org/10.1016/j.artmed.2024.102861>
19. Eke CI, Shuib L. The role of explainability and transparency in fostering trust in AI healthcare systems: a systematic literature review, open issues and potential solutions. *Neural Computing and Applications*. 2025;37:1999-2034. <https://doi.org/10.1007/s00521-024-10868-x>
20. Singh RP, Hom GL, Abramoff MD, Campbell JP, Chiang MF. Current Challenges and Barriers to Real-World Artificial Intelligence Adoption for the Healthcare System, Provider, and the Patient. *Translational Vision Science & Technology*. 2020;9(2):45. <https://doi.org/10.1167/tvst.9.2.45>
21. Carroll NW, Jones A, Burkard T, Lulias C, Severson K, Posa T. Improving risk stratification using AI and social determinants of health. *American Journal of Managed Care*. 2022;28(11):582-587. <https://doi.org/10.37765/ajmc.2022.89261>
22. Nong P, Adler-Milstein J. Socially situated risk: challenges and strategies for implementing algorithmic risk scoring for care management. *JAMIA Open*. 2021;4(3):ooab076. <https://doi.org/10.1093/jamiaopen/ooab076>
23. Uslu A, Stausberg J. Value of the Electronic Medical Record for Hospital Care: Update From the Literature. *Journal of Medical Internet Research*. 2021;23(12):e26323. <https://doi.org/10.2196/26323>
24. Sharda M, Sharma S, Raikar S, et al. The Role of Machine Learning in Predicting Hospital Readmissions Among General Internal Medicine Patients: A Systematic Review. *Cureus*. 2025;17(5):e84761. <https://doi.org/10.7759/cureus.84761>
25. Arbaein T. Preventable Hospitalization and Primary Healthcare Visits Among Hypertensive Patients in Makkah City. *Healthcare*. 2025;13(23):3039. <https://doi.org/10.3390/healthcare13233039>
26. Rocha JVM, Santana R, Tello JE. Hospitalization for ambulatory care sensitive conditions: What conditions make inter-country comparisons possible? *Health Policy OPEN*. 2021;2:100030. <https://doi.org/10.1016/j.hpopen.2021.100030>
27. Eldesoky HAM, AlThubaity D, Shalby AYM, Mohammed FA. Solicitude toward artificial intelligence among health care providers and its relation to their patient's safety culture in Saudi

- Arabia. BMC Health Services Research. 2025;25:889. <https://doi.org/10.1186/s12913-025-13001-3>
28. Fischer GS, da Rosa Righi R, de Oliveira Ramos G, da Costa CA, Rodrigues JJPC. ElHealth: Using Internet of Things and data prediction for elastic management of human resources in smart hospitals. Engineering Applications of Artificial Intelligence. 2020;87:103285. <https://doi.org/10.1016/j.engappai.2019.103285>
29. Suleiman AK, Ming LC. Transforming healthcare: Saudi Arabia's vision 2030 healthcare model. Journal of Pharmaceutical Policy and Practice. 2025;18(1):2449051. <https://doi.org/10.1080/20523211.2024.2449051>
30. Daghistani T, Alshammari R. Comparison of statistical logistic regression and random forest machine learning techniques in predicting diabetes. Journal of Advances in Information Technology. 2020;11(2):78-83. <https://doi.org/10.12720/jait.11.2.78-83>