Explainable AI for Population Health Analytics: A Conceptual Framework for Multicultural Healthcare Decision Support Systems

Osama Sohaib

Department of Statistics and Business Analytics, College of Business and Economics

United Arab Emirates University, UAE

osohaib@uaeu.ac.ae

Abstract

The United Arab Emirates faces a substantial burden of non-communicable diseases. Diabetes prevalence exceeds global averages, and cardiovascular disease accounts for a significant share of national mortality. What distinguishes the UAE's public health challenge from many other contexts is not merely the scale of disease burden but the composition of the population. The country's population includes substantial expatriate communities whose cultural backgrounds, health beliefs, and healthcare-seeking patterns vary considerably. A one-size-fits-all approach to public health intervention is therefore inadequate. This paper introduces CAXAI, a conceptual framework that combines health data, population-stratified analytics, and Explainable Artificial Intelligence methods to generate risk factor insights tailored to specific cultural groups. The framework aims to support evidence-based health policy and resource allocation consistent with the UAE Vision 2030.

Keywords: *Explainable AI; Multicultural Public Health; Non-communicable Diseases; Health Disparities; UAE;*

1. INTRODUCTION

Non-communicable diseases represent the dominant cause of mortality globally. The World Health Organization estimates they account for roughly three quarters of deaths worldwide, making them a defining public health crisis of our time (WHO, 2022). The UAE is no exception to this trend, though the burden here manifests with particular intensity. Diabetes affects a substantially larger share of the adult population than the global average, while cardiovascular disease is responsible for a disproportionate fraction of deaths recorded nationally (IDF, 2021; MoHAP, 2020).

The public health response to this burden is complicated by a fundamental demographic reality. The country's population is predominantly composed of expatriates. Over 200 nationalities are represented, creating what may be one of the most culturally diverse societies on earth (FCSC, 2021). This diversity shapes health in multiple ways. Healthcare utilisation patterns differ markedly across cultural groups. So do the underlying disease risks, the health beliefs that guide decisions about seeking care, and the ways that material conditions and cultural adaptation interact to produce health outcomes (Shehab et al., 2023; Sulaiman et al., 2017).

Artificial intelligence has demonstrated considerable promise in healthcare, showing performance advantages over traditional methods in disease prediction and risk stratification (Topol, 2019). Yet most AI systems operate as black boxes. This opacity limits their usefulness, particularly in public health contexts where understanding why a population faces elevated risk matters as much as accurately predicting that risk. The lack of transparency also raises legitimate ethical concerns. When AI models are deployed in diverse populations, unexamined biases embedded in training data can amplify existing health inequities (Rajkomar et al., 2018). Explainable AI offers a different path forward. By making model reasoning transparent and interpretable, XAI techniques such as SHAP and LIME enable decision-makers to understand not just what an AI system predicts but why (Arrieta et al., 2020). Applied to public health in multicultural settings, XAI can reveal how risk factors differ across population subgroups, pointing toward interventions that actually address the specific conditions shaping health outcomes for particular communities (Panigutti et al., 2021).

This paper presents CAXAI: Culturally-Aware Explainable AI for public health. It is a conceptual framework designed to apply XAI methods to understanding differential NCD risk across the UAE's diverse population. Four research questions guide the work. First, how can XAI reveal population-specific risk factor patterns in multicultural settings? Second, what interpretable insights might guide culturally sensitive public health strategies? Third, can explainable AI actually support evidence-based policymaking in demographically complex populations? Fourth, can XAI insights be embedded within intelligent health information systems to support resource allocation decisions?

2. BACKGROUND

2.1 Non-communicable Disease Burden and Health Disparities in the UAE

The UAE's NCD profile reflects the country's rapid urbanisation. Dietary patterns have shifted substantially from traditional foods toward processed alternatives. Physical activity has declined due to sedentary work patterns and climate conditions that make outdoor activity challenging during much of the year (Alhyas et al., 2011). What emerges is a distinctive epidemiological picture. Health outcomes vary considerably between Emirati citizens and expatriate populations. Emiratis show elevated rates of diabetes and obesity. South Asian expatriates, by contrast, experience higher cardiovascular disease risk despite often having lower obesity levels (Shehab et al., 2023; Sulaiman et al., 2017). These disparities cannot be attributed to genetic factors alone. Cultural health beliefs play a role. Language barriers affect healthcare access and quality. Socioeconomic conditions differ substantially across groups. All of these factors interact in ways that produce unequal health outcomes (Al-Maskari et al., 2010).

2.2 AI and Explainable AI in Healthcare Settings

Machine learning has evolved well beyond basic statistical analysis. It now provides genuine predictive capability that can inform clinical and public health decisions (Khoury et al., 2016; Beam & Kohane, 2018). The challenge, however, is that most algorithms function as black boxes. Accountability becomes difficult when no one can explain how a system reached its conclusion. Clinicians worry about hidden biases in the underlying training data. Patients and communities worry about whether they can actually trust the system's recommendations (Chen et al., 2020). Explainable AI addresses this problem directly. Methods like SHAP, LIME, and counterfactual explanations make model reasoning transparent without sacrificing accuracy (Lundberg & Lee, 2017; Ribeiro et al., 2016; Wachter et al., 2017). In clinical practice, this matters. XAI has helped clinicians understand why sepsis risk models flag particular patients as high-risk (Komorowski et al., 2018). In oncology, XAI explanations support more transparent diagnostic reasoning (Shifa et al., 2025). Yet in multicultural public health specifically, where understanding population-level disparities is the goal, XAI applications remain underdeveloped. The literature on XAI tends to focus on individual prediction rather than group-level pattern detection. Few frameworks address how to use XAI to uncover differential risk factors across demographically distinct communities.

2.3 Research Gaps

Current literature reveals four critical gaps: (1) most XAI healthcare studies focus on homogeneous populations, lacking frameworks for differential risk factor analysis across demographic subgroups; (2) existing approaches inadequately integrate cultural health beliefs and social determinants; (3) systematic mechanisms for translating XAI insights into culturally tailored public health policies are absent; and (4) comprehensive empirical validation of XAI approaches within real-world, diverse healthcare systems is lacking (Arrieta et al., 2020; Panigutti et al., 2021). This study will address these gaps.

3. THE CAXAI CONCEPTUAL FRAMEWORK

3.1 Framework Overview

The CAXAI framework integrates five interconnected layers: (1) Multi-source Data Integration, (2) Population Stratification, (3) Ensemble Predictive Modelling, (4) Multi-method XAI Analysis, and (5) a Policy Translation Interface. Figure 1 illustrates the complete architecture.

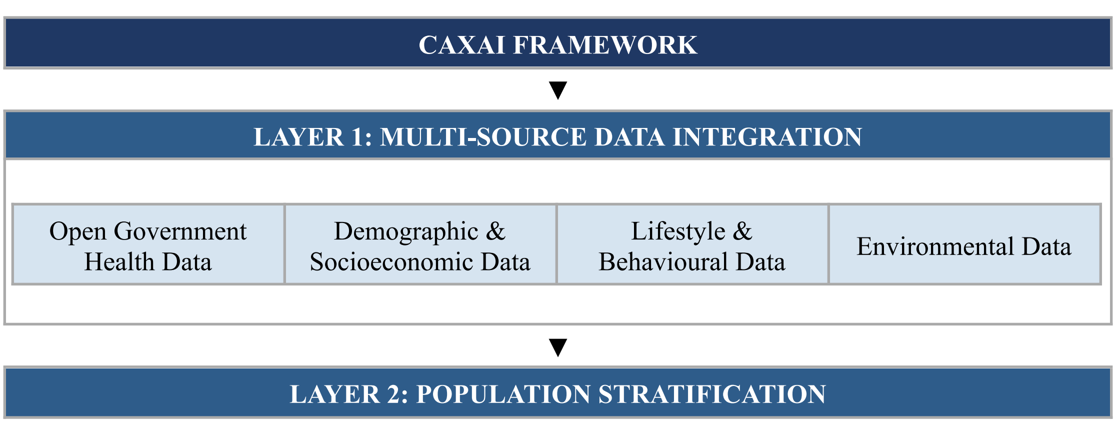

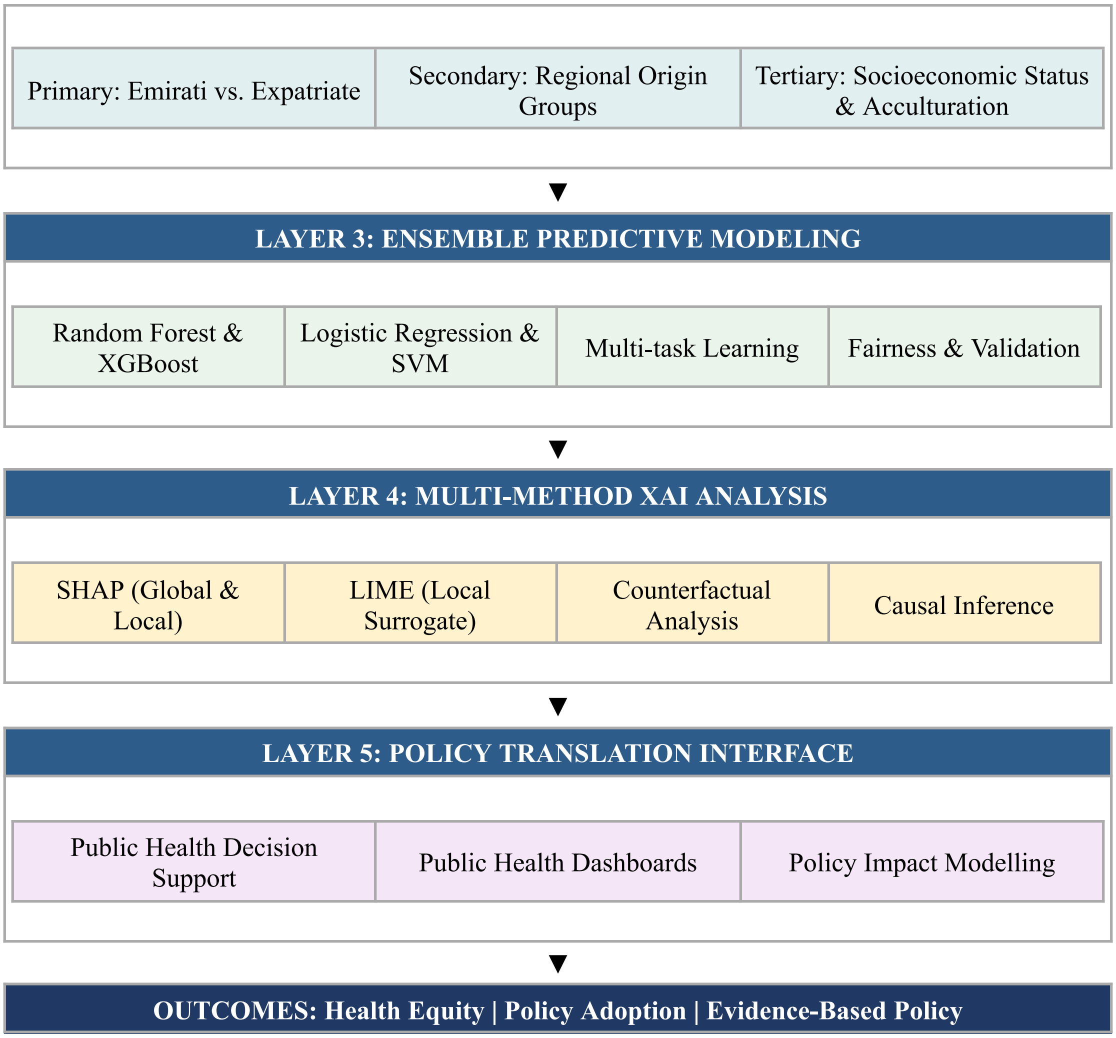

Figure 1: CAXAI Framework Architecture for Explainable AI in UAE Multicultural Public Health

3.2 Layer 1: Multi-Source Data Integration

The framework's first phase involves compiling data exclusively from publicly accessible open data sources within the UAE. Core health and demographic information will be obtained from the UAE Open Data Portal, which centralizes official datasets from government entities such as the Ministry of Health and Prevention (MoHAP) and the Federal Competitiveness and Statistics Centre (FCSC). These resources offer population-level insights into NCD prevalence, mortality, and demographic indicators, including nationality, age, gender, socioeconomic status, and employment. Additional data on behavioural and lifestyle risk factors are drawn from publicly released findings of the UAE National Health Survey and WHO STEPS surveillance reports for the UAE and GCC region. Environmental data, including air quality indices, temperature patterns, and built environment characteristics, are accessed through open datasets from the UAE National Centre of Meteorology and municipal open data portals. Because all the data comes from aggregated, anonymized, and publicly available sources at the population level, there's no involvement of individual patient records, clinical diagnoses, or lab results. This means privacy risks are very low but even so, all data handling will still follow UAE Federal Law No. 45 of 2021 on Personal Data Protection to make sure everything is done responsibly and ethically.

3.3 Layer 2: Population Stratification

Given the UAE's demographic complexity, a three-level stratification strategy will be employed. Stratification Levels: Primary stratification differentiates Emirati citizens from expatriates. Secondary stratification groups expatriates by region of origin. Tertiary stratification considers socioeconomic factors. Stratification Purpose: To capture genetic, cultural, and socioeconomic influences on healthcare interactions and health outcomes. Stratification Techniques: Both supervised demographic categorisation and unsupervised clustering techniques, including Gaussian Mixture Models and hierarchical clustering, are used (Fraley and Raftery, 2002).

3.4 Layer 3: Ensemble Predictive Modelling

The modelling layer employs an ensemble approach combining tree-based methods (Random Forest, XGBoost) for capturing non-linear relationships within aggregated population data, logistic regression and Support Vector Machines (SVM) for interpretable classification of group-level risk patterns, and linear methods for identifying key relationships across demographic strata. This combination is particularly suited to the structured, aggregated nature of open government datasets. A multi-task learning architecture simultaneously models multiple NCD outcomes (diabetes, hypertension, cardiovascular disease, obesity), illuminating shared and distinct risk factor patterns across demographic groups (Caruana, 1997). Rather than transfer learning from individual-level clinical datasets, the framework draws on publicly available international NCD risk factor reports and WHO regional datasets for contextual benchmarking and model calibration. Model fairness is assessed throughout using demographic parity, equalised odds, and cross-group calibration metrics to ensure equitable outputs across population subgroups (Barocas et al., 2019).

3.5 Layer 4: Multi-Method XAI Analysis

This layer constitutes the analytical core of CAXAI, deploying three complementary XAI methodologies. Global explanations use SHAP (SHapley Additive Explanations) to quantify overall feature importance within each demographic subgroup, revealing which factors most strongly drive NCD risk across the population (Lundberg & Lee, 2017). Partial Dependence Plots and permutation importance provide complementary insights into nonlinear relationships. Local explanations employ LIME (Local Interpretable Model-agnostic Explanations) to generate individual-level risk factor profiles for high-risk persons, while anchor explanations identify minimal feature sets sufficient to determine risk classifications (Ribeiro et al., 2016, 2018). Counterfactual analysis via the DiCE framework generates actionable, diverse "what-if" scenarios, thus identifying minimal changes to personal characteristics that would alter risk predictions, directly informing intervention design (Mothilal et al., 2020). Causal discovery algorithms (PC algorithm, Fast Causal Inference) complement XAI outputs by distinguishing correlation from causation, while formal mediation analysis identifies pathways through which cultural and socioeconomic factors shape health disparities (VanderWeele, 2015).

3.6 Layer 5: Policy Translation Interface

The framework concludes with a policy translation interface converting XAI outputs into stakeholder-specific actions. Public health decision support tools provide population-level risk insights with interpretable explanations for health authorities and programme managers, supporting evidence-based health planning and prioritisation. Public health dashboards present population-level risk factor visualisations and disparity analyses for policymakers. A population-level intervention recommendation engine generates culturally tailored strategies based on group-level risk patterns across demographic communities. Policy impact modelling enables predictive analysis of proposed health interventions' effects on population health outcomes and equity measures, directly supporting UAE Vision 2030 health objectives.

4. DISCUSSION

CAXAI makes three categories of contribution. Methodologically, it advances XAI by introducing population-stratified explainability and cultural context integration capabilities absent from existing frameworks (Arrieta et al., 2020). In practical terms, this approach will yield novel insights into how non-communicable disease risks vary across the diverse population groups in the UAE, helping to measure health inequalities and pinpoint actionable targets for public health interventions. From a policy perspective, using transparent AI to generate evidence—rather than relying on assumptions or sparse observational data—will support more accurate distribution of resources, culturally sensitive health initiatives, and policies aimed at advancing health equity.

Several challenges must be acknowledged. Ensuring the completeness, currency, and granularity of open government datasets presents a key challenge, as publicly available data may be aggregated at levels that limit subgroup analysis precision. Cultural context quantification presents methodological complexity; collaboration with anthropologists and mixed-methods approaches are proposed as mitigations. Ensemble approaches combining interpretable and complex models are needed to balance interpretability and performance. Algorithmic fairness requires continuous monitoring and bias mitigation (Barocas et al., 2019), while stakeholder buy-in necessitates engagement and policy relevance.

While designed for the UAE, CAXAI principles are generalisable to other multicultural societies confronting similar NCD challenges, including Gulf Cooperation Council countries, immigrant-destination nations, and diverse urban populations globally. The framework's explicit attention to cultural context, fairness, and interpretability sets new standards for ethical AI development in global health settings.

5. CONCLUSION

This paper has presented CAXAI, a conceptual framework applying Explainable AI to precision public health in the UAE's culturally diverse population. By integrating multi-source open data, population-stratified predictive modelling, and complementary XAI techniques within a structured policy translation interface, the framework addresses the critical gap between AI predictions and policy-actionable, culturally relevant insights. In a country where the vast majority of the population comprises expatriates from over 200 nationalities, and where the NCD burden is severe, such interpretable precision approaches are not merely desirable but essential. Future work will focus on empirical validation through open data analysis, pilot policy applications, and iterative refinement of the policy translation interface to support meaningful health equity gains.

Acknowledgement

This study was funded by the UAEU University under Grant No. 12B103.

REFERENCES

Al-Maskari, F., El-Sadig, M., & Nagelkerke, N. (2010). Assessment of the direct medical costs of diabetes mellitus and its complications in the United Arab Emirates. BMC Public Health, 10(1), 679

Alhyas, L., McKay, A., Balasanthiran, A., & Majeed, A. (2011). Prevalences of overweight, obesity, hyperglycaemia, hypertension and dyslipidaemia in the Gulf: Systematic review. JRSM Short Reports, 2(7), 55. https://doi.org/10.1258/shorts.2011.011019

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82–115. https://doi.org/10.1016/j.inffus.2019.12.012

Barocas, S., Hardt, M., & Narayanan, A. (2023). *Fairness and machine learning: Limitations and opportunities*. MIT press.

Beam, A. L., & Kohane, I. S. (2018). Big data and machine learning in health care. JAMA, 319(13), 1317–1318. https://doi.org/10.1001/jama.2017.18391

Caruana, R. (1997). Multitask learning. Machine Learning, 28(1), 41–75. https://doi.org/10.1023/A:1007379606734

Chen, I. Y., Joshi, S., Ghassemi, M., & Ranganath, R. (2020). Treating health disparities with AI. Nature Medicine, 26(1), 16–17. https://doi.org/10.1038/s41591-019-0649-2

Fraley, C., & Raftery, A. E. (2002). Model-based clustering, discriminant analysis, and density estimation. Journal of the American Statistical Association, 97(458), 611–631. https://doi.org/10.1198/016214502760047131

Holzinger, A., Langs, G., Denk, H., Zatloukal, K., & Müller, H. (2019). Causability and explainability of artificial intelligence in medicine. WIREs Data Mining and Knowledge Discovery, 9(4), e1312. https://doi.org/10.1002/widm.1312

International Diabetes Federation. (2021). IDF Diabetes Atlas (10th ed.). IDF.

Khoury, M. J., Iademarco, M. F., & Riley, W. T. (2016). Precision public health for the era of precision medicine. American Journal of Preventive Medicine, 50(3), 398–401. https://doi.org/10.1016/j.amepre.2015.08.031

Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C., & Faisal, A. A. (2018). The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. Nature Medicine, 24(11), 1716–1720. https://doi.org/10.1038/s41591-018-0213-5

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774.

Mothilal, R. K., Sharma, A., & Tan, C. (2020). Explaining machine learning classifiers through diverse counterfactual explanations. FAccT '20 Proceedings, 607–617. https://doi.org/10.1145/3351095.3372850

Panigutti, C., Perotti, A., Panisson, A., Bajardi, P., & Pedreschi, D. (2021). FairLens: Auditing black-box clinical decision support systems. *Information Processing & Management*, *58*(5), 102657.

Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2018). Ensuring fairness in machine learning to advance health equity. Annals of Internal Medicine, 169(12), 866–872. https://doi.org/10.7326/M18-1990

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. KDD '16 Proceedings, 1135–1144. https://doi.org/10.1145/2939672.2939778

Ribeiro, M. T., Singh, S., & Guestrin, C. (2018). Anchors: High-precision model-agnostic explanations. AAAI Conference on Artificial Intelligence, 1527–1535.

Shehab, A., Bakir, S., Sabbour, H., Elnour, A. A., Mahmeed, W. A., Salam, A. M., & Kholy, D. E. (2023). Prevalence of cardiovascular risk factors and 10-years risk for coronary heart disease in the United Arab Emirates. *Current Diabetes Reviews*, *19*(3), 38-48.

Shifa, N., Saleh, M., Akbari, Y., & Al Maadeed, S. (2025). A review of explainable AI techniques and their evaluation in mammography for breast cancer screening. *Clinical Imaging*, *123*, 110492.

Sulaiman, N., Elbadawi, S., Hussein, A., Abusnana, S., Madani, A., Mairghani, M., ... & Peeters, A. (2017). Prevalence of overweight and obesity in United Arab Emirates Expatriates: the UAE national diabetes and lifestyle study. *Diabetology & metabolic syndrome*, *9*(1), 88.

Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. Nature Medicine, 25(1), 44–56. https://doi.org/10.1038/s41591-018-0300-7

UAE Federal Competitiveness and Statistics Centre. (2021). Population statistics. UAE Government.

UAE Ministry of Health and Prevention. (2020). Health Statistics Annual Report. MoHAP.

VanderWeele, T. J. (2015). Explanation in causal inference: Methods for mediation and interaction. Oxford University Press.

Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. Harvard Journal of Law & Technology, 31(2), 841–887.

World Health Organization. (2022). Health literacy development for the prevention and control of non-communicable diseases, Volume 2.

Author Bio

Dr. Osama Sohaib is an Associate Professor of Business Analytics in the Department of Statistics and Business Analytics at the UAE University. He holds a Ph.D. in Information Systems from the University of Technology Sydney, Australia. His expertise spans business analytics, applied machine learning, digital transformation, and information systems. His research focuses on leveraging data-driven technologies to support organizational decision-making, innovation, and digital transformation.