(10) Publication No: **IN202641010439 A1**
(22) Date of Filing: 01-02-2026
(43) Publication Date: 13-02-2026
Journal No: 07/2026

(19) Intellectual Property Office, India

(12) INDIAN PATENT APPLICATION

(54) Title: **Secured breath sensor array framework for non-invasive disease diagnosis using generative artificial intelligence**

(51) International Classification:
A61B 5/00, A61B 5/083, A61B 5/097, A61B 5/08 A61B 5/1455

(21) Application No: **202641010439**

(31) Priority Document No: --

(32) Priority Date: --

(86) International Application No: --
Filing Date: --

(87) International Publication No: --

(61) Patent of Addition to Application No: --
Filing Date: --

(62) Divisional to Application No: --
Filing Date: --

(71) Name of Applicant(s):
Vellore Institute Of Technology,

(72) Name of Inventor(s):
1. Dr. Yokesh Babu Sundaresan
2. Dr. Kumaresan P
3. Oshim Pathan
4. Rahul Yadav
5. Aashish Kumar Mahato

(57) Abstract:

ABSTRACT SECURED BREATH SENSOR ARRAY FRAMEWORK FOR NON-INVASIVE DISEASE DIAGNOSIS USING GENERATIVE ARTIFICIAL INTELLIGENCE The present disclosure provides a sensor array system including a hybrid sensor array (1, 2, 3, 5) having a plurality of gas sensors of different sensing modalities configured to detect volatile organic compounds in exhaled breath, including at least one metal oxide semiconductor sensor (1, 2), at least one non-dispersive infrared sensor (3), and at least one mass-sensitive sensor (5). The system further includes physiological sensors (8, 9, 10, 11) configured to acquire vital sign data from a patient and an edge computing platform (13) communicatively coupled to the hybrid sensor array and physiological sensors (8, 9, 10, 11). The edge computing platform (13) fuses outputs from the hybrid sensor array and physiological sensors (8, 9, 10, 11) to generate fused sensor outputs and processes them using embedded machine learning algorithms to generate a disease risk index. (FIG. 2)

FORM-2

THE PATENTS ACT, 1970

(39 of 1970)

&

THE PATENTS RULES, 2003

COMPLETE

SPECIFICATION

(See section 10 and rule 13)

TITLE OF INVENTION

SECURED BREATH SENSOR ARRAY FRAMEWORK FOR NON-INVASIVE DISEASE DIAGNOSIS USING GENERATIVE ARTIFICIAL INTELLIGENCE

APPLICANT NAME AND ADDRESS

Vellore Institute of Technology

An Indian University having address as:

Vellore Campus, Vellore - 632 014

Tamil Nadu, India

THE FOLLOWING SPECIFICATION PARTICULARLY DESCRIBES THE INVENTION AND THE MANNER IN WHICH IT IS TO BE PERFORMED

CROSS-REFERENCE TO RELATED APPLICATIONS AND PRIORITY

[0001] The present application does not claim priority from any patent application.

PREAMBLE

[0002] The following specification particularly describes the invention and the manner in which it is to be performed.

FIELD OF INVENTION

[0003] The present disclosure relates to biomedical sensing systems integrating wearable electronics, respiratory monitoring, breath analysis, and Internet of Things connectivity with machine learning, and more particularly to a secured breath sensor array framework combining hybrid gas sensor arrays with physiological vital sign sensors and edge-based generative artificial intelligence for continuous non-invasive disease diagnosis and telemedicine workflow support.

BACKGROUND

[0004] The following text of background art is provided for purposes of understanding and is not an admission that any cited document or technique forms part of the prior art under applicable law.

[0005] Wearable health monitoring devices have emerged as tools for continuous physiological surveillance, enabling real-time acquisition of vital parameters such as heart rate, blood pressure, oxygen saturation, and body temperature. Concurrently, breath analysis has demonstrated potential for non-invasive detection of volatile organic compounds (VOCs) that may serve as biomarkers for metabolic, respiratory, hepatic, renal, and cardiac conditions. US 11,045,111 B1 discloses a real-time breath analyzer for detecting VOCs using sensor arrays with graphite-polymer composites and pattern recognition algorithms for disease identification, though the system focuses on diagnostic workflows rather than continuous wearable monitoring with multimodal vitals fusion. US 2022/0240808 A1 describes a disposable electrochemical wearable sensor for breath biochemistry monitoring using porous substrates with differential electrode designs, but does not address hybrid metal oxide semiconductor, non-dispersive infrared, and quartz crystal microbalance sensor arrays integrated with physiological monitoring in an edge machine learning pipeline. CN 112394172 A targets acetone monitoring for diabetic patients using portable devices, though the

approach concentrates on single biomarker detection rather than heterogeneous sensor arrays fused with vital sign data. US 2007/0062255 A1 and US 2015/0177224 A1 describe handheld breath collection and acetone monitoring devices respectively, without provisions for continuous wearable operation, edge-based inference, or clinician workflow integration. BR 112018070768 A8 discloses breath analysis for VOC detection but does not specify hybrid sensor configurations with drift compensation or integrated clinical dashboards.

[0006] Existing wearable health systems exhibit several technical limitations. Many commercial devices rely on single-modality sensing without integration of biochemical breath markers, constraining their utility for early-stage disease detection. Standalone breath analysis devices remain largely confined to laboratory settings or single-disease diagnostic applications without real-time, continuous, or remote monitoring capabilities. Current wearables predominantly employ threshold-based alerts rather than sophisticated multi-parameter disease risk scoring using machine learning or deep learning methods. Physiological and biochemical data streams are often siloed across disparate systems, complicating telemedicine integration and reducing predictive accuracy. Existing solutions frequently lack robust data security, privacy protection, and compliance with healthcare regulations when managing sensitive patient data across cloud platforms. Machine learning in biomedical domains suffers from limited access to large, clinically diverse datasets, restricting model generalization for rare or early-stage conditions.

[0007] It has been appreciated that a system is needed that overcomes one or more of these problems.

[0008] The present invention addresses the above shortcomings of the prior art. However, the invention is entirely different from the prior art in terms of novelty and technological advancements.

SUMMARY

[0009] In a first aspect, a sensor array system is provided. The sensor array system comprises a hybrid sensor array including a first metal oxide semiconductor sensor and a second metal oxide semiconductor sensor configured to detect ammonia, carbon dioxide, and benzene associated with kidney disorders, liver disorders, and lung disorders.

[0010] In a second aspect, a system is provided. The system comprises the sensor array system of the first aspect, wherein the system processes fused outputs of the sensor array using embedded machine learning algorithms deployed on an edge computing platform.

[0011] In a third aspect, a system is provided. The system comprises the system of the second aspect, wherein the system integrates a large language model configured to summarize historical patient health data and real-time patient health data into interpretable clinical reports.

[0012] In a fourth aspect, a system is provided. The system comprises the system of the first aspect, the system of the second aspect, and the system of the third aspect, wherein the system provides a secure web platform with role-based dashboards for patients, doctors, and administrators.

[0013] In a fifth aspect, a method for non-invasive disease diagnosis is provided. The method comprises collecting breath analysis data from a patient using a hybrid sensor array comprising a first metal oxide semiconductor sensor, a second metal oxide semiconductor sensor, a non-dispersive infrared sensor, a multi-gas sensor, a quartz crystal microbalance sensor, and a chemocapacitor sensor, wherein the first metal oxide semiconductor sensor and the second metal oxide semiconductor sensor are configured to detect ammonia, carbon dioxide, and benzene, the non-dispersive infrared sensor is configured to measure carbon dioxide concentration, the multi-gas sensor is configured to detect ethanol and volatile organic compounds, the quartz crystal microbalance sensor is configured to detect acetone, and the chemocapacitor sensor is configured to detect volatile organic compounds associated with infections and cancer markers.

[0014] In a sixth aspect, a non-transitory computer-readable medium is provided. The non-transitory computer-readable medium stores instructions that, when executed by a processor, cause the processor to receive fused sensor outputs from a hybrid sensor array comprising metal oxide semiconductor sensors, a non-dispersive infrared sensor, a multi-gas sensor, a quartz crystal microbalance sensor, and a chemocapacitor sensor, process the fused sensor outputs using machine learning algorithms comprising a random forest model, a support vector machine model, and a long short-term memory recurrent neural network, generate a disease risk index and a confidence score based on the processed fused sensor outputs, invoke a large language model to generate an

interpretable clinical report summarizing the disease risk index, the confidence score, and historical patient health data, and transmit the interpretable clinical report to a role-based dashboard interface for display to an authorized user.

[0015] The foregoing general description of the illustrative embodiments and the following detailed description thereof are merely exemplary aspects of the teachings of this disclosure and are not restrictive.

BRIEF DESCRIPTION OF FIGURES

[0016] The drawings appended herein illustrate exemplary embodiments of the present invention and are provided to enhance understanding of the inventive features disclosed. These figures, when read in conjunction with the detailed description, depict the operational flow, component architecture, and other mechanisms.

[0017] FIG. 1 illustrates a block diagram of a secured breath sensor array framework for non-invasive disease diagnosis, according to aspects of the present disclosure.

[0018] FIG. 2 illustrates a top view of a wearable sensor array system configured for health monitoring and breath analysis, according to an embodiment.

[0019] FIG. 3 illustrates a multi-model health assessment pipeline for processing patient health data, according to aspects of the present disclosure.

[0020] FIG. 4 illustrates a flowchart for a machine learning processing workflow deployed on an edge computing platform, according to an embodiment.

[0021] FIG. 5 illustrates a flowchart for a clinical report generation process using a large language model, according to aspects of the present disclosure.

[0022] FIG. 6 illustrates a flowchart for an AI prediction generation workflow within an administrator dashboard interface, according to an embodiment.

[0023] FIG. 7 illustrates a flowchart for a patient ID generation algorithm, according to aspects of the present disclosure.

[0024] FIG. 8 illustrates a system diagram of a breath sensor array framework for non-invasive disease diagnosis, according to an embodiment.

[0025] FIG. 9 illustrates a layered system architecture diagram depicting a multi-tier dashboard interface structure, according to aspects of the present disclosure.

[0026] Furthermore, in terms of the construction of the system, one or more components of the device may have been represented in the figures by conventional

symbols, and the figures may show only those specific details that are pertinent to understanding the embodiments of the present invention so as not to obscure the figures with details that will be readily apparent to those of ordinary skill in the art having benefit of the description herein.

DETAILED DESCRIPTION

[0027] The following detailed description is intended to provide an in-depth understanding of the invention and its various components, configurations, and operational aspects. While specific embodiments are described with reference to the accompanying drawings, the invention is not limited to the particular forms disclosed herein. The scope of the invention shall be interpreted broadly and in accordance with the claims appended hereto.

[0028] The terms "comprises," "comprising," "includes," "including," and similar expressions used in this description are intended to be open-ended and should be interpreted to include, but not be limited to, the mentioned elements. Unless explicitly stated otherwise, singular terms may include their plural equivalents and vice versa, depending on the context in which they appear.

[0029] The present invention will now be described in detail with reference to the accompanying drawings illustrating exemplary embodiments.

[0030] The present disclosure relates to a secured breath sensor array framework for non-invasive disease diagnosis using generative artificial intelligence. The framework provides a multimodal wearable health monitoring platform that integrates physiological monitoring capabilities with breath analysis functionality and machine learning-based disease risk prediction.

[0031] The framework addresses limitations in existing health monitoring systems by combining multiple sensing modalities within a unified platform. Conventional wearable devices typically focus on single-modality sensing of physiological parameters such as heart rate, blood pressure, oxygen saturation, and body temperature. Such conventional approaches lack integration with biochemical markers that provide additional diagnostic information for early-stage disease detection.

[0032] Breath analysis provides a non-invasive approach to detecting volatile organic compounds in exhaled breath. These volatile organic compounds serve as biomarkers indicative of metabolic, renal, hepatic, respiratory, cardiac, and infectious

conditions. The framework combines breath analysis with physiological monitoring to enable comprehensive health assessment through a single wearable platform.

[0033] The framework incorporates edge computing capabilities for real-time processing of sensor data using embedded machine learning algorithms. The edge computing approach enables on-device inference for disease risk prediction without requiring continuous cloud connectivity. Cloud-based processing supplements edge computing by providing aggregate data analytics and personalized model refinement based on longitudinal patient data.

[0034] The framework further incorporates generative artificial intelligence capabilities for synthesizing clinical narratives from patient health data. A large language model processes historical and real-time health data to generate interpretable summaries for healthcare providers. The combination of machine learning-based risk scoring and generative artificial intelligence-based summarization provides clinicians with both quantitative risk assessments and qualitative health narratives.

[0035] The framework provides a telehealth platform with role-based access control for patients, healthcare providers, and administrators. Each user role accesses functionality and data appropriate to the designated role within the health monitoring ecosystem. The platform enables real-time data visualization, historical trend analysis, alert management, remote prescription, and appointment scheduling through secure web-based interfaces.

[0036] Data security and privacy protection are implemented throughout the framework through encrypted data transmission, secure authentication, and role-based access control mechanisms. The framework is designed for compliance with healthcare data regulations governing the handling of sensitive patient information across cloud platforms and multiple user roles.

[0037] FIG. 1 illustrates a block diagram of the secured breath sensor array framework according to various embodiments. The framework comprises multiple interconnected layers that facilitate data acquisition, cloud management, user interaction, AI-based prediction, clinical decision support, and patient outcome delivery.

[0038] The data acquisition layer includes a patient ID generation module 101 configured to generate unique patient identifiers for secure data linkage. The patient ID

generation module 101 generates a unique patient identifier using a country code extraction, region suffix determination, and check digit computation algorithm. The patient ID generation module 101 employs Verhoeff, Luhn, or ISO/IEC 7064 check digit algorithms for error detection in patient identifiers. The patient ID generation module 101 includes concurrency protection with atomic insert operations guarded by unique database index constraints to prevent ID collisions during simultaneous identifier generation requests.

[0039] The data acquisition layer further includes a vital data collection 102 configured for continuous acquisition of physiological parameters such as heart rate, blood pressure, oxygen saturation, and temperature. A breath analysis data collection module 103 captures volatile organic compounds and related biomarkers from exhaled breath. The vital data collection 102 and the breath analysis data collection module 103 operate in parallel to provide multimodal health data for subsequent processing and analysis.

[0040] The cloud data management layer receives data from the data acquisition layer and includes an encrypt and upload module 201 configured to encrypt all collected sensor data before secure transmission to cloud storage for patient privacy protection. A step 202 provides longitudinal storage of patient health records over time. The system stores patient health records in a Supabase cloud database for longitudinal tracking of health status and trends. A data synchronization module 203 automatically synchronizes health data across all users and authorized clinicians for real-time updates. The data synchronization module 203 ensures that current health information is available for analysis and consultation across the platform.

[0041] The web application user management layer provides role-based access through a patient dashboard 301 that displays personal health trends, alerts, risk scores, appointments, and medication information to patients. A step 302 provides a doctor dashboard for reviewing patient lists, detailed risk summaries, predictive analytics, vitals graphs, and disease analysis trends. The doctor dashboard supports remote prescription issuance and appointment scheduling functionality. An administrator dashboard 303 enables management of user access, patient reports, and data integrity across the system.

[0042] The system implements secure authentication with optional two-step verification and session management for protected area access. The system implements role-based access control mechanisms compliant with HIPAA and GDPR healthcare data regulations. Each user role accesses functionality and data appropriate to the designated role within the health monitoring ecosystem.

[0043] The AI and ML prediction layer includes an ML prediction engine 401 that processes real-time data to predict disease risks and flag health changes. The ML prediction engine 401 integrates with a Gemini API large language model to generate health summaries and risk assessments from patient data. The ML prediction engine 401 implements retry logic with bounded retries and backoff for LLM API calls to handle transient failures. A step 402 provides an LLM engine for summarization that processes historical and new data to synthesize insights and generate clinical narratives.

[0044] The system implements timestamp-based cache validation comparing cache timestamps with latest vitals and breath data timestamps to determine prediction freshness. The system stores prediction results in a database cache with input timestamps for deterministic application-layer reuse until upstream data changes. This caching approach minimizes latency and computational cost while preserving data freshness through timestamp comparison.

[0045] The doctor decision support layer includes an ML predictions display module 501 that presents machine learning predictions and LLM summaries in an actionable format during clinical consultations. The ML predictions display module 501 includes ML disease prediction display showing confidence percentages and individual disease risk percentages for conditions including lung infection, hypotension, and inflammation. A graphs and alerts module 502 generates visualizations including graphs, trends, and alert displays to assist clinicians in interpreting health trajectories. A diagnosis assistance module 503 combines real-time analytics and comprehensive summaries to support informed and timely clinical diagnoses.

[0046] The patient outcomes layer includes a continuous health monitoring module 601 that provides ongoing non-invasive health surveillance. A personalized insights and alerts module 602 delivers automated personalized health feedback and early warnings to improve patient engagement. The system utilizes event-tiered alert

protocols with dynamic thresholds personalized to patient baselines and medical history to minimize false positives. The system implements adaptive patient-specific baselines that are refined over time through cloud-assisted learning.

[0047] An AI-assisted decision support module 603 enables clinicians to leverage AI-driven support for rapid and reliable decision-making. A single platform convenience module 604 consolidates all processes including data collection, analysis, consultation, alerts, and historical review within a unified digital platform.

[0048] The system targets diagnostic accuracy greater than 90% across multiple disease categories using advanced explainable machine learning models. The system enables detection of chronic diseases including diabetes, chronic kidney disease, hepatic disorders, respiratory ailments, cardiac abnormalities, and infections.

[0049] The system is designed for scalable cloud deployment using cloud data lakes and secure wireless connectivity. The system employs modular hardware and software architectures that accommodate future sensor expansions and analytic algorithm upgrades. The system architecture is prepared for clinical translation via pathways incorporating device calibration, user acceptability testing, and clinical validation aligned with ISO 13485 and FDA standards.

[0050] FIG. 2 illustrates a top view of a wearable sensor array system according to various embodiments. FIG. 8 illustrates a system diagram of the breath sensor array framework according to various embodiments. The wearable sensor array system comprises a hybrid sensor array mounted on a breadboard 4 that serves as the central mounting platform for integrating multiple sensor modules and an edge computing platform.

[0051] The hybrid sensor array includes a MQ-135 metal oxide semiconductor (MOX) sensor 1 configured to detect ammonia, carbon dioxide, and benzene associated with kidney, liver, and lung disorders. The MQ-135 metal oxide semiconductor (MOX) sensor 1 provides multi-analyte detection capability for identifying volatile compounds that serve as biomarkers for metabolic dysfunction across multiple organ systems. The MQ-135 metal oxide semiconductor (MOX) sensor 1 operates through changes in electrical resistance when target gas molecules interact with the metal oxide sensing layer.

[0052] A MQ-137 metal oxide semiconductor (MOX) sensor 2 is positioned adjacent to the MQ-135 metal oxide semiconductor (MOX) sensor 1 on the breadboard 4. The MQ-137 metal oxide semiconductor (MOX) sensor 2 is configured to detect ammonia as a biomarker for kidney and liver dysfunction. A breath inlet tube MQ-137 connects to the MQ-137 metal oxide semiconductor (MOX) sensor 2 for directing exhaled breath samples to the sensor array. The breath inlet tube MQ-137 provides a controlled pathway for breath sample collection and delivery to the gas sensing elements. The breath analysis data collection module 103 receives breath samples through the breath inlet tube MQ-137 connected to the hybrid sensor array.

[0053] An MH-Z19 non-dispersive infrared (NDIR) sensor 3 is mounted on the breadboard 4 and is configured to measure carbon dioxide concentration in exhaled breath for respiratory and metabolic disorder detection. The MH-Z19 non-dispersive infrared (NDIR) sensor 3 employs infrared absorption spectroscopy to quantify CO2 levels in exhaled breath with high specificity. The MH-Z19 non-dispersive infrared (NDIR) sensor 3 provides accurate CO2 measurements that correlate with respiratory function and metabolic state.

[0054] A MiCS-5524 multi-gas sensor 6 is included in the hybrid sensor array and is configured to detect ethanol and volatile organic compounds linked to liver dysfunction and lung diseases. The MiCS-5524 multi-gas sensor 6 provides broad-spectrum VOC detection capability for identifying multiple breath biomarkers associated with hepatic and pulmonary conditions.

[0055] A quartz crystal microbalance (QCM) sensor 5 is mounted on the breadboard 4 and is configured to detect acetone as a biomarker for diabetes. The quartz crystal microbalance (QCM) sensor 5 operates through frequency shifts induced by mass changes on the crystal surface when acetone molecules adsorb to the sensing layer. The quartz crystal microbalance (QCM) sensor 5 provides sensitive detection of acetone concentrations that correlate with diabetic ketoacidosis and metabolic dysregulation.

[0056] A chemocapacitor sensor module 7 is positioned on the breadboard 4 and is configured to detect VOCs associated with early-stage infections and cancer markers. A Chemocapacitor sensor for infection/cancer marker 12 provides additional detection capability for volatile compounds indicative of infectious processes and malignant

conditions. The chemocapacitor sensor module 7 and the Chemocapacitor sensor for infection/cancer marker 12 operate through capacitance changes induced by VOC adsorption on dielectric sensing layers.

[0057] The hybrid sensor array includes on-body drift compensation and temperature compensation mechanisms integrated within a single wearable platform. The drift compensation mechanisms correct for sensor baseline shifts that occur during extended continuous monitoring periods. The temperature compensation mechanisms adjust sensor readings based on ambient and body temperature variations to maintain measurement accuracy across operating conditions.

[0058] For physiological monitoring, the wearable sensor array system incorporates ECG electrode pads 8 and ECG electrode pads 9 for cardiac signal acquisition. The ECG electrode pads 8 and the ECG electrode pads 9 attach to the patient's body for collecting ECG signals as part of continuous cardiac monitoring. The vital data collection 102 collects ECG signals through the ECG electrode pads 8 and the ECG electrode pads 9 attached to the patient's body.

[0059] A temperature sensor 10 is provided for body temperature measurement as part of the integrated vital signs monitoring system. The temperature sensor 10 measures body temperature continuously to detect fever and temperature-related physiological changes.

[0060] A blood pressure sensor 11 is included for blood pressure monitoring. The blood pressure sensor 11 measures systolic and diastolic blood pressure values as part of continuous vital monitoring. The blood pressure sensor 11 provides periodic or continuous blood pressure readings for cardiovascular health assessment.

[0061] The vital data collection 102 collects SpO2 oxygen saturation measurements as part of continuous physiological monitoring. An SpO2 pulse oximetry sensor module provides oxygen saturation readings through optical measurement of blood oxygen levels.

[0062] A Raspberry Pi 4 edge computing platform 13 is mounted on the breadboard 4 and serves as the processing unit for the sensor array. The Raspberry Pi 4 edge computing platform 13 receives data from all connected sensors and executes embedded machine learning algorithms for real-time disease risk prediction and data fusion. The Raspberry Pi 4 edge computing platform 13 includes I2C pins, analog

inputs, and SPI connections for interfacing with the various sensors. The Raspberry Pi 4 edge computing platform 13 includes a power input, a power switch, and a MicroSD card slot for storage. The Raspberry Pi 4 edge computing platform 13 processes fused outputs from the sensor array using embedded machine learning algorithms to generate disease risk indices and confidence scores for clinical decision support.

[0063] The arrangement of components on the breadboard 4 facilitates integration of multiple sensing modalities within a compact form factor suitable for wearable health monitoring applications. The system is designed for high wearability with optimized battery life for continuous monitoring in clinical and homecare environments. The wearable design incorporates power management strategies to extend operational duration while maintaining continuous data acquisition from all sensor modalities.

[0064] FIG. 3 illustrates a multi-model health assessment pipeline according to various embodiments. The multi-model health assessment pipeline processes patient health data through parallel analytical pathways to generate unified health reports for clinical decision support.

[0065] The multi-model health assessment pipeline begins with patient health data input, which undergoes preprocessing and data validation. The pipeline receives data from multiple sources including wearable device data from the Raspberry Pi 4 edge computing platform 13, electronic health records, and breath analysis sensor data from the breath analysis data collection module 103.

[0066] Following preprocessing, the validated data enters a parallel processing stage where two distinct analytical pathways operate simultaneously. A first pathway involves Gemini AI analysis, which proceeds through Gemini AI validation and subsequently generates a textual risk assessment. A second pathway involves ML model prediction, which proceeds through ML model validation, and the ML model generates probability scores for various health conditions.

[0067] The outputs from both analytical pathways converge at a results combination stage. The combined results proceed through a model comparison phase, followed by cross-model validation to check agreement and divergence between the AI-generated textual assessment and the ML-generated probability scores. The model

comparison and cross-model validation stages detect failure modes and quantify reliability of the hybrid system outputs.

[0068] A confidence scoring module evaluates the reliability of the combined predictions using agreement patterns, calibration metrics, and input quality assessments. The confidence scoring supports safer triage of outputs for clinician review by quantifying the consistency between parallel analytical pathways.

[0069] The pipeline culminates in unified health report generation that produces comprehensive documentation integrating both the textual risk assessment from the Gemini AI pathway and the quantitative probability scores from the ML model pathway. The unified health reports combine text narratives with numeric risk scores, rationale explanations, and caveats for clinician review. A doctor dashboard display presents the analyzed information to healthcare providers for clinical decision-making.

[0070] FIG. 4 illustrates a flowchart for a machine learning processing workflow deployed on an edge computing platform according to various embodiments. The flowchart depicts the sequence of operations for processing fused sensor outputs and generating disease risk predictions within a patient dashboard interface 100.

[0071] The process begins with a step 104 of receiving fused outputs from the hybrid sensor array. The fused outputs comprise data collected from multiple sensors including the vital data collection 102 and breath analysis sensors. Following data reception, the process proceeds to a step 106 where the fused outputs are processed using embedded machine learning algorithms.

[0072] The Raspberry Pi 4 edge computing platform 13 executes real-time inference using Random Forest algorithms for disease risk prediction. The Raspberry Pi 4 edge computing platform 13 executes real-time inference using Support Vector Machine algorithms for disease risk prediction. The Raspberry Pi 4 edge computing platform 13 executes real-time inference using Deep Neural Networks with attention mechanisms for disease risk prediction. The Raspberry Pi 4 edge computing platform 13 executes real-time inference using Long Short-Term Memory (LSTM) recurrent neural networks for temporal pattern detection. The LSTM recurrent neural networks analyze sequential sensor data to identify temporal disease patterns that evolve over monitoring periods.

[0073] The workflow advances to a step 108 where the machine learning algorithms are deployed on the Raspberry Pi edge platform. After deployment, the process reaches a decision point at a step 110 where the system determines whether the processing was successful. If the processing is successful, the workflow proceeds to step 108 where the system generates a disease risk index and confidence score. The ML prediction engine 401 generates a disease risk index on a scale of 1-10 with an associated confidence score. The ML prediction engine 401 employs a model ensemble and consensus mechanism combining multiple model outputs for reliable disease risk assessment. These outputs assist clinical decision-making by providing quantified risk assessments based on the multimodal sensor data.

[0074] If the processing is not successful, the workflow diverts to step 110 where the system logs the error and initiates a retry of the processing operation. The error handling mechanism ensures robustness in the edge computing pipeline by capturing failure events and attempting recovery through reprocessing.

[0075] The system implements a FastAPI-based ML service workflow with CORS middleware for request validation and route handling. External client requests pass through CORS middleware to validate origin, and invalid origins are rejected with a CORS validation failed outcome. Valid requests proceed to route handling for matched paths including predict, health, and root endpoints.

[0076] The ML prediction engine 401 performs input validation checking schema, types, ranges, and required fields before feature extraction. Validation errors are surfaced as input validation error responses, while valid inputs trigger feature extraction and cleaning operations.

[0077] The ML prediction engine 401 performs feature extraction and cleaning using Pandas and NumPy libraries to coerce, impute, scale, encode, and assemble model-ready feature vectors. The system implements feature categorization splitting variables into numeric features such as heart rate, SpO2, temperature, and blood pressure, and categorical features such as device type and sampling mode for appropriate transformation processing.

[0078] The system applies numeric normalization including z-score and min-max scaling to stabilize training and inference across sensor ranges. The numeric

normalization aligns value ranges across different sensor modalities to prevent any single sensor from dominating the feature space.

[0079] The system applies categorical encoding via one-hot, ordinal, target, frequency, or binary encoders depending on cardinality and model requirements. The system employs hash-based encoding for high-cardinality VOC types and ECG signal labels to control dimensionality and handle unseen categories at inference time.

[0080] The ML prediction engine 401 loads serialized scikit-learn model artifacts via Joblib for model prediction operations. At startup or on first call, model loading retrieves a serialized scikit-learn artifact from the file system. The loaded estimator is used for model prediction, and failures are reported as model prediction error responses.

[0081] The ML prediction engine 401 implements health check endpoints for readiness and liveness checks enabling orchestrators and load balancers to gate traffic when the model and dependencies are ready. The health endpoint runs health check logic including model-in-memory verification, filesystem and joblib availability, and dependency checks.

[0082] The system performs shape validation and feature list verification to confirm expected feature count and names matching the model's training contract. The shape validation prevents column mismatch failures in production by enforcing deterministic column order that the serialized estimator expects.

[0083] The system performs probability array shape checking after model inference to validate output dimensions before response formatting. The probability shape checking validates K-class softmax length to catch model or contract regressions.

[0084] The system performs output bounds validation to sanity-check prediction scores and ensure they meet expected ranges. The output bounds validation guards against corrupted artifacts or overflow conditions by verifying that classification scores sum appropriately and regression outputs fall within expected bounds.

[0085] The system employs synthetic dataset generation using state-of-the-art data augmentation approaches to enable AI model training despite limited clinical datasets. The synthetic data generation improves model robustness, privacy protection, and generalization for rare disease conditions.

[0086] FIG. 5 illustrates a flowchart for clinical report generation using a large language model according to various embodiments. The process operates within a doctor dashboard interface 200 to transform raw patient health data and machine learning predictions into clinician-friendly narrative summaries.

[0087] The process begins with a step 202 where the system receives patient health data from the machine learning system. The process proceeds to a step 204 where the system retrieves historical patient health records from the database. The historical records provide longitudinal context for the current health assessment.

[0088] The process moves to a step 206 where a large language model is integrated for data summarization. The large language model processes the combined current and historical health data to generate interpretable clinical narratives.

[0089] The process advances to a step 208 where a decision is made to determine whether the summarization was successful. If the summarization is successful, the process proceeds to a step 210 where the system generates an interpretable clinical report for the clinician. The clinical report includes narrative summaries highlighting abnormal deviations and longitudinal health trajectory analysis.

[0090] If the summarization is not successful, the process returns to a retry operation at step 202 where the large language model summarization is retried with error handling. The retry mechanism incorporates bounded retries and backoff to handle transient faults or rate limits in LLM API calls.

[0091] FIG. 6 illustrates a flowchart for an AI prediction generation workflow according to various embodiments. The workflow operates within an administrator dashboard interface 300 and implements timestamp-based cache validation to minimize latency and computational cost while preserving data freshness.

[0092] The workflow begins with a step 302 where the system inputs and validates a patient ID to ensure the user exists and is authorized. The workflow proceeds to a step 304 where the system fetches the latest vitals and breath data from the database.

[0093] The workflow moves to a step 306 where the system retrieves the prediction cache and associated timestamps to determine cache validity. At step 306, a decision is made to check if the cache is valid by comparing the cache timestamp with the newest data timestamps.

[0094] If the cache is valid, the workflow proceeds to a step 308 where the system returns the cached structured prediction to avoid recomputation and reduce latency. The cached prediction includes the patient identifier, input timestamps, model version, narrative summary, risk scores with confidences, and alert information.

[0095] If the cache is not valid, the workflow moves to a step 310 where the system formats the health data for AI analysis. The formatting builds a structured, normalized feature and context representation for consistent LLM prompting.

[0096] The workflow proceeds to a step 312 where the system generates a summary and risk assessment via the Gemini API. The summary generation and risk assessment generation employ bounded retries and backoff to withstand transient faults or rate limits.

[0097] Upon successful generation, the workflow moves to a step 314 where the system caches the results in the database and returns the structured prediction. The caching persists the successful prediction artifact including summary and risk vector with input timestamps for instant responses on subsequent requests until upstream data changes.

[0098] FIG. 7 illustrates a flowchart for a patient ID generation algorithm according to various embodiments. The algorithm generates unique patient identifiers with collision prevention and error detection capabilities.

[0099] The process begins with a step 400 where the system validates the country name and extracts the country code from the input data. The country code extraction uses ISO country code standards. The system determines a region or site suffix and creates a country profile for applying country-specific rules including prefix, encoding, and reserved ranges.

[0100] The process proceeds to a step 402 where the system queries existing IDs for a deduplication check to ensure uniqueness of the identifier being generated. If the database operation fails, the system surfaces a database error.

[0101] At a step 404, the system determines whether a potential duplicate is detected. If a potential duplicate is detected due to race or concurrency conditions, the process moves to a step 406 where the system triggers concurrency protection and retry mechanisms. The concurrency protection includes retry operations or switching to a sequence generator to handle collision scenarios.

[0102] If no potential duplicate is detected, the process proceeds to a step 408 where the system assembles a candidate ID with a base prefix derived from the validated country information. The ID assembly includes building a candidate suffix with sequence or entropy components and performing level checks to avoid orphan or digits-only identifiers.

[0103] The process continues to a step 410 where the system computes a check digit and validates the final ID format. The check digit computation uses Verhoeff, Luhn, or ISO/IEC 7064 algorithms for error detection when identifiers are typed or transmitted manually. The validation confirms prefix, length, allowed characters, and check digit correctness.

[0104] The process moves to a step 412 where the system persists the identifier via an atomic insert operation and returns the unique patient ID. The atomic insert is guarded by a unique database index constraint. If the insert fails because another thread raced to the same ID, the workflow re-enters the increment and regeneration loop until success. The database-level uniqueness constraint serves as the single source of truth for collision prevention.

[0105] The real-time health data synchronization module implements data type routing splitting flows into vitals monitoring data and breath analysis data for insertion into separate tables. The typed destination enforcement simplifies lineage tracking, quality rules, and performance tuning per data stream.

[0106] The system implements fault tolerance on database write failures including error logging, dead letter queue routing, connection retry, and transaction rollback. The fault tolerance mechanisms reduce data loss and support idempotent design for reliable pipeline operation. Successful inserts trigger real-time dashboard updates and cache invalidation for predictions to maintain consistency with the source of truth.

[0107] The system performs data aggregation for historical analysis including temporal grouping, statistical aggregation, anomaly detection, and time-series organization. The historical aggregation supports cohort analysis and time window computations for population health insights.

[0108] The system performs correlation analysis and trend identification for clinical insights and monitoring model stability. The correlation analysis identifies

relationships between sensor readings and health outcomes over longitudinal monitoring periods.

[0109] The system generates data quality metrics and predictive feature generation to enrich ML inputs and monitor pipeline health. The data quality metrics include freshness, volume, schema compliance, and distribution monitoring to detect data drift and quality degradation.

[0110] FIG. 9 illustrates a layered system architecture diagram depicting the multi-tier dashboard interface structure according to various embodiments. The system is organized into three distinct layers, each corresponding to a specific user role within the telehealth ecosystem. The three layers are interconnected through a hierarchical data flow structure with bidirectional communication pathways between the layers.

[0111] Layer 1 is designated as the patient dashboard interface and is positioned at the top of the architecture. The patient dashboard interface provides functionality tailored for patient users. The patient dashboard interface includes a personal health trends module that displays longitudinal health data and trajectory information to patients. A wearable device interface module enables patients to view data collected from the sensor array components and manage device connectivity settings. An alert notification module delivers automated notifications regarding health status changes, appointment reminders, and medication schedules. A medication tracking panel enables patients to monitor prescribed medications, dosages, and adherence schedules. A risk score visualization panel presents disease risk indices generated by the machine learning prediction engine in a graphical format that patients can interpret. An appointment management panel enables patients to view scheduled appointments, request new appointments, and manage appointment preferences.

[0112] Layer 2 is designated as the doctor dashboard interface and is positioned in the middle tier of the architecture. The doctor dashboard interface provides functionality designed for healthcare providers. A patient list panel displays all patients assigned to the healthcare provider with summary health status indicators. A vitals analysis data display module presents physiological parameter data including heart rate, blood pressure, oxygen saturation, and temperature measurements collected by the sensor array. A breath analysis display module presents volatile organic compound

measurements and breath biomarker data collected by the hybrid sensor array including ammonia, carbon dioxide, acetone, ethanol, and benzene concentrations.

[0113] The doctor dashboard interface further incorporates a machine learning disease prediction panel that displays confidence percentages for predicted health conditions. The machine learning disease prediction panel presents individual disease risk percentages for conditions including lung infection, hypotension, and inflammation alongside an overall prediction confidence score. An artificial intelligence health prediction and risk assessment panel presents narrative summaries generated by the large language model that synthesize patient health data into interpretable clinical reports. A medical history panel provides access to longitudinal patient health records and historical trend data. An appointment scheduling panel enables healthcare providers to schedule patient consultations, manage appointment availability, and coordinate follow-up visits.

[0114] Layer 3 is designated as the administrator dashboard interface and is positioned at the bottom tier of the architecture. The administrator dashboard interface provides functionality for system administrators. A user management panel enables administrators to create, modify, and deactivate user accounts across all user roles within the platform. The user management panel supports role assignment and access permission configuration. A data integrity search and management panel enables administrators to search patient records, verify data consistency, and resolve data quality issues across the platform. A patient report statistics panel generates aggregate reports on patient populations, system usage metrics, and platform performance indicators for system oversight and quality assurance purposes.

[0115] The sensor array components collect physiological and breath biomarker data from patients through continuous monitoring. The physiological sensors including the ECG electrode pads, temperature sensor, blood pressure sensor, and SpO2 pulse oximetry sensor acquire vital sign measurements. The breath analysis sensors including the MOX sensors, NDIR sensor, QCM sensor, multi-gas sensor, and chemocapacitor sensors acquire volatile organic compound measurements from exhaled breath. The collected sensor data is transmitted to the edge computing platform for machine learning processing. The edge computing platform executes embedded machine

learning algorithms to perform feature extraction, data fusion, and disease risk prediction on the multimodal sensor data.

[0116] The processed data flows through the cloud data management layer for secure storage and synchronization. The encrypt and upload module encrypts all sensor data and machine learning outputs before transmission to cloud storage. The longitudinal storage component maintains patient health records over time to enable historical trend analysis and personalized baseline refinement. The data synchronization module distributes updated health information across all authorized users and clinicians to ensure that current data is available for analysis and consultation through the role-based dashboards.

[0117] The AI and ML prediction layer generates disease risk indices and large language model-based summaries from the processed sensor data. The machine learning prediction engine produces quantitative disease risk scores with associated confidence values for multiple health conditions. The large language model engine processes historical and current health data to generate narrative summaries that synthesize complex multimodal data into interpretable clinical reports. The generated predictions and summaries are presented through the role-based dashboards according to user access permissions.

[0118] The doctor decision support layer combines machine learning predictions with visualizations to assist clinical diagnosis. The machine learning predictions display module presents disease risk indices and confidence scores in an actionable format during clinical consultations. The graphs and alerts module generates trend visualizations, comparative charts, and alert indicators that enable healthcare providers to interpret patient health trajectories. The diagnosis assistance module integrates quantitative risk assessments with qualitative narrative summaries to support informed clinical decision-making.

[0119] The patient outcomes layer provides continuous monitoring, personalized alerts, and AI-assisted decision support through a unified platform. The continuous health monitoring module maintains ongoing non-invasive surveillance of patient health status through the wearable sensor array. The personalized insights and alerts module delivers automated health feedback and early warning notifications based on dynamic thresholds personalized to individual patient baselines. The AI-assisted

decision support module enables healthcare providers to leverage machine learning predictions and large language model summaries for rapid clinical assessment. The single platform convenience module consolidates all health monitoring, analysis, consultation, and communication functions within a unified digital interface accessible to patients, healthcare providers, and administrators according to their designated roles.

[0120] There are several technical advantages of the secured breath sensor array framework of the present disclosure. First, the integration of a hybrid sensor array comprising metal oxide semiconductor sensors, non-dispersive infrared sensors, quartz crystal microbalance sensors, and chemocapacitor sensors within a single wearable platform enables simultaneous detection of multiple breath biomarkers including ammonia, carbon dioxide, benzene, acetone, ethanol, and volatile organic compounds, thereby providing comprehensive non-invasive disease screening across multiple organ systems without requiring separate diagnostic devices for each target analyte. Second, the deployment of embedded machine learning algorithms on an edge computing platform enables real-time disease risk prediction at the point of data collection, reducing latency in clinical decision support and eliminating dependence on continuous network connectivity for generating diagnostic assessments. Third, the fusion of breath analysis data with physiological vital sign data including electrocardiogram signals, oxygen saturation, temperature, and blood pressure measurements creates a multimodal health assessment capability that improves diagnostic accuracy by capturing interdependencies between biochemical and physiological health indicators that would not be detectable from either data modality in isolation. Fourth, the integration of a large language model for generating interpretable clinical reports transforms complex multimodal sensor outputs and machine learning predictions into human-readable narrative summaries, reducing cognitive burden on healthcare providers and facilitating rapid comprehension of patient health status during clinical consultations. Fifth, the role-based telehealth platform with differentiated dashboards for patients, healthcare providers, and administrators ensures that each user category accesses functionality and data appropriate to their designated role while maintaining data compartmentalization and compliance with healthcare data privacy regulations through encrypted transmission, authentication, and role-based access control mechanisms. Sixth, the timestamp-based cache validation mechanism minimizes computational cost and

latency while preserving data freshness, and the synthetic data generation approaches address data scarcity challenges in biomedical machine learning to enable robust model training across diverse patient populations including rare disease conditions. These technical advantages collectively contribute to a comprehensive solution for continuous non-invasive health monitoring that bridges the gap between breath-based diagnostics and practical telemedicine workflows.

[0121] It will be appreciated that the above-detailed description, along with FIGs, is provided to illustrate the architecture and operation of exemplary embodiments of the present invention. Various modifications and enhancements can be made without departing from the scope of the invention. All such variations, as would be recognized by those skilled in the art, are intended to be within the scope of the present disclosure, which is defined by the claims that follow.

[0122] It will be appreciated by persons skilled in the art that while the invention has been described with reference to specific embodiments and accompanying drawings, numerous modifications, substitutions, variations, and equivalents are possible without departing from the spirit or scope of the present invention. The use of singular terms such as “a,” “an,” and “the” is not intended to limit the disclosed elements to a single instance unless explicitly stated, and such terms should be interpreted to include plural forms as applicable. Likewise, the terms “comprises,” “comprising,” “includes,” “including,” “has,” “having,” and their grammatical variants are intended to be open-ended and non-limiting, and should be interpreted as meaning “including but not limited to.” Any enumerated listing of components, features, or elements does not imply that the items are mutually exclusive unless expressly stated.

[0123] It is further understood that where features, characteristics, or elements are described in connection with a Markush group or a disjunctive phrase (e.g., “A or B”), the invention includes all possible combinations and sub-combinations thereof unless explicitly excluded. Thus, a statement referring to “at least one of A, B, and C” should be interpreted to mean any one or more of A, B, and C, individually or in any combination. The language used in this specification is selected primarily for clarity and illustrative purposes and should not be construed to limit the inventive scope unless specifically recited in the claims. Accordingly, the foregoing description of

embodiments should be regarded as illustrative and not restrictive, with the scope of the invention being defined solely by the appended claims and their legal equivalents

[0124] It will be recognized that various features, elements, and combinations thereof described herein may be desirably adapted for alternative implementations or other applications. Many such modifications, substitutions, enhancements, or equivalents may become apparent to those skilled in the art upon reading this disclosure and are considered to fall within the scope and spirit of the present invention. The invention is not limited to the precise configurations or exemplary embodiments set forth in this specification, and various alternatives may be used without departing from the intended objectives. The claims, and not the detailed description, shall define the legal scope of protection afforded by the present disclosure.

[0125] In interpreting the specification and claims, all terms should be construed in the broadest reasonable manner consistent with the context and the understanding of those skilled in the art. The use of terms such as “comprises,” “comprising,” “includes,” “including,” or “has” should be interpreted to be non-exclusive and not limited to the stated elements alone. Moreover, references to “one embodiment,” “an embodiment,” or “in certain embodiments” are not meant to imply that such features are required or exclusive, and any feature described in connection with one embodiment may be used in combination with other features or embodiments unless clearly prohibited or contextually inconsistent.

[0126] It should be further appreciated that reference throughout the specification to features, operations, or components using singular terms should not be construed as excluding a plurality thereof unless the context expressly indicates otherwise. Similarly, features described in the context of grouped or listed elements (e.g., A, B, and C) should be understood to encompass individual, multiple, or all combinations of said elements. Where ranges are given, all intermediate values and subranges are understood to be disclosed as if specifically recited. The scope of the invention should therefore be construed as inclusive of all such combinations and logical extensions as would be appreciated by a person of ordinary skill in the relevant technical field

[0127] In another embodiment, the above disclosure is a description of the invention and is not intended to limit the scope of the invention. Other variations and modifications of the above-described embodiment shall be apparent to those skilled in

the art and are intended to fall within the scope of the invention as defined in the following claims.

Dated This 31st day of January 2026

(Kuldeep Singh)

Authorized Agent for the Applicant,

Indian Patent Agent Regn No. IN/PA-4358

CLAIMS

WE CLAIM:

1. A sensor array system comprising:

a hybrid sensor array comprising a plurality of gas sensors of different sensing modalities configured to detect volatile organic compounds in exhaled breath, wherein the plurality of gas sensors includes at least one metal oxide semiconductor sensor (1, 2), at least one non-dispersive infrared sensor (3), and at least one mass-sensitive sensor (5);

a plurality of physiological sensors (8, 9, 10, 11) configured to acquire vital sign data from a patient;

an edge computing platform (13) communicatively coupled to the hybrid sensor array and the plurality of physiological sensors (8, 9, 10, 11);

wherein the edge computing platform (13) is configured to fuse outputs from the hybrid sensor array and the plurality of physiological sensors (8, 9, 10, 11) to generate fused sensor outputs; and

wherein the edge computing platform (13) is configured to process the fused sensor outputs using embedded machine learning algorithms to generate a disease risk index.

2. The sensor array system of claim 1, wherein the at least one metal oxide semiconductor sensor (1, 2) is configured to detect ammonia, carbon dioxide, and benzene associated with kidney disorders, liver disorders, and lung disorders.

3. The sensor array system of claim 1 or 2, wherein the at least one non-dispersive infrared sensor (3) is configured to measure carbon dioxide concentration in exhaled breath for respiratory disorder detection and metabolic disorder detection.

4. The sensor array system of any of claims 1 to 3, wherein the at least one mass-sensitive sensor comprises a quartz crystal microbalance sensor (5) configured to detect acetone as a biomarker for diabetes.

5. The sensor array system of any of claims 1 to 4, further comprising a chemocapacitor sensor (7, 12) configured to detect volatile organic compounds associated with early-stage infections and cancer markers.

6. The sensor array system of any of claims 1 to 5, wherein the edge computing platform (13) is further configured to generate a confidence score associated with the

disease risk index, wherein the confidence score is configured to provide a reliability metric for clinical decision support.

7. A health monitoring system comprising:

the sensor array system of any of claims 1 to 6;

a cloud data management layer comprising an encrypt and upload module (201) configured to receive encrypted data from the edge computing platform (13) and a data synchronization module (203) configured to store and synchronize patient health records;

a machine learning prediction engine (401) configured to generate disease risk assessments from the fused sensor outputs;

a large language model configured to generate interpretable clinical reports from patient health data; and

a plurality of role-based dashboard interfaces comprising a patient dashboard interface (100), a doctor dashboard interface (200), and an administrator dashboard interface (300), wherein the plurality of role-based dashboard interfaces are configured to provide differentiated access to patient health data and disease risk assessments based on user role.

8. The health monitoring system of claim 7, further comprising a cache validation mechanism configured to compare a cache timestamp with timestamps of latest vitals data and breath analysis data to determine prediction freshness, wherein the cache validation mechanism is configured to return a cached structured prediction when the cache is valid and to trigger generation of a new prediction when the cache is invalid.

9. A computer-implemented method for generating disease risk assessments, the method comprising:

receiving (104) fused sensor outputs from a hybrid sensor array comprising a plurality of gas sensors (1, 2, 3, 5) configured to detect breath biomarkers and a plurality of physiological sensors (8, 9, 10, 11) configured to acquire vital sign data;

processing (106) the fused sensor outputs using embedded machine learning algorithms deployed on an edge computing platform (13); and

generating (108) a disease risk index based on the processed fused sensor outputs.

10. A non-transitory computer-readable medium storing instructions that, when executed by a processor of an edge computing platform (13), cause the processor to:

receive fused sensor outputs from a hybrid sensor array comprising a plurality of gas sensors (1, 2, 3, 5) configured to detect breath biomarkers and a plurality of physiological sensors (8, 9, 10, 11) configured to acquire vital sign data;

process the fused sensor outputs using machine learning algorithms comprising at least one of a random forest model, a support vector machine model, and a long short-term memory recurrent neural network; and

generate a disease risk index based on the processed fused sensor outputs.

Dated This 31st day of January 2026

Kuldeep

(Kuldeep Singh)

Authorized Agent for the Applicant,

Indian Patent Agent Regn No. IN/PA-4358

ABSTRACT

SECURED BREATH SENSOR ARRAY FRAMEWORK FOR NON-INVASIVE DISEASE DIAGNOSIS USING GENERATIVE ARTIFICIAL INTELLIGENCE

The present disclosure provides a sensor array system including a hybrid sensor array (1, 2, 3, 5) having a plurality of gas sensors of different sensing modalities configured to detect volatile organic compounds in exhaled breath, including at least one metal oxide semiconductor sensor (1, 2), at least one non-dispersive infrared sensor (3), and at least one mass-sensitive sensor (5). The system further includes physiological sensors (8, 9, 10, 11) configured to acquire vital sign data from a patient and an edge computing platform (13) communicatively coupled to the hybrid sensor array and physiological sensors (8, 9, 10, 11). The edge computing platform (13) fuses outputs from the hybrid sensor array and physiological sensors (8, 9, 10, 11) to generate fused sensor outputs and processes them using embedded machine learning algorithms to generate a disease risk index.

(FIG. 2)

Dated This 31st day of January 2026

Kuldeep

(Kuldeep Singh)

Authorized Agent for the Applicant,

Indian Patent Agent Regn No. IN/PA-4358