# Title

# High-Throughput AI-Enhanced FT-IR Spectroscopic Diagnostic System for Humoral Primary Immunodeficiency and Non-Communicable Disease Screening

## Description

## Field of the Invention

The present invention relates to the field of in-vitro diagnostic systems and spectroscopic analysis. In particular, it concerns a high-throughput Fourier-transform infrared (FT-IR) spectroscopic platform enhanced with artificial intelligence for screening humoral primary immunodeficiencies and assessing risk indicators of non-communicable diseases from micro-volume biological samples.

## Background of the Invention

Primary immunodeficiencies (PIDs) affecting the humoral immune system (e.g., X-linked agammaglobulinemia, common variable immunodeficiency) often go undiagnosed for years due to the lack of accessible screening methods. Current newborn screening methods (such as the T-cell receptor excision circle (TREC) PCR assay) detect cellular PIDs but fail to identify over 90% of humoral PIDs, which rely on measuring antibody levels. Existing quantitative immunoglobulin tests (e.g., nephelometry, ELISA, Luminex-based assays) require reagents, cold-chain logistics, and skilled technicians in laboratory settings. In low-resource and remote settings, these requirements are difficult to meet, resulting in many immunodeficiencies remaining undetected until severe symptoms manifest.

Non-communicable diseases (NCDs) such as diabetes, dyslipidemia, malnutrition, sepsis risk, and kidney disease represent a parallel public health burden, especially in low- and middle-income regions. Diagnosing or screening for these conditions typically involves multiple separate assays (for glucose, cholesterol, albumin, etc.), which are time-consuming and resource-intensive. There is an unmet need for a reagent-free, rapid, and multi-analyte screening test that can simultaneously assess immunoglobulin levels and key metabolic health indicators with minimal sample volume.

Spectroscopic analysis techniques, including infrared spectroscopy, have been investigated for clinical diagnostics because they can capture broad biochemical information from a small sample without reagents. Prior art has demonstrated the feasibility of using spectroscopy with algorithmic processing to estimate analyte concentrations. For example, near-infrared (NIR) spectrometry combined with machine learning has been used to non-invasively predict blood analytes, and Raman spectroscopy with AI has been applied for real-time material quality control. However, no existing solution provides a point-of-care screening platform that fulfills the WHO A.S.S.U.R.E.D. criteria (Affordable, Sensitive, Specific, User-friendly, Rapid, Robust, Equipment-free, and Deliverable) for immunoglobulin and multi-analyte testing. In particular, no current system can produce a multi-analyte panel from a single ≤5 µL blood sample in under one minute without reagents. The present invention addresses these gaps by integrating FT-IR spectroscopy with an AI-driven analysis pipeline and innovative sample handling, enabling population-scale screening for humoral PIDs and NCD risk in both well-equipped laboratories and resource-limited clinics.

# Summary of the Invention

In light of the aforementioned needs and deficiencies in the prior art, the present invention provides a comprehensive AI-enhanced FT-IR spectroscopic diagnostic system for rapid, multi-analyte screening of immune function and metabolic health using only a micro-volume blood sample. The system uniquely integrates a mid-infrared spectrometer with specialized sample handling hardware and a suite of algorithmic modules to translate a raw infrared spectrum into clinically relevant outputs in real time. The following key aspects characterize the invention and distinguish it from prior systems:

First, the invention features a micro-volume ATR (Attenuated Total Reflection) sampling tile (602) that requires only about 3 µL of serum or plasma. This disposable tile includes a micro-patterned surface (e.g. a hydrophobic perimeter and hydrophilic center) which self-distributes a droplet into a thin, uniform film of 2–5 µm thickness. This ensures a consistent optical pathlength for the FT-IR measurement and obviates the need for complex manual spreading. The tile is packaged in a desiccated foil pouch (610) with a barcode (612) for quality control and traceability, maintaining spectral baseline stability over time. In one embodiment, the tiles can be automatically conveyed in sequence across the spectrometer's ATR sensor, with on-board cleaning (e.g. an isopropanol wipe and UV sterilization) between samples to enable high-throughput operation (on the order of up to 1000 tests per hour).

Secondly, the system incorporates an age-adaptive chemometric analysis engine (706). Unlike prior spectroscopic diagnostic methods that apply a one-size-fits-all model (often calibrated only on adults), this invention maintains a library of calibration models stratified by age group. An age-band selection module automatically determines the patient's age (for example, infant, child, or adult band) and selects the appropriate multivariate calibration vector or machine learning model for that cohort before making predictions. This age-specific model selection ensures that normal developmental differences in the spectral signatures (such as protein levels in infants vs. adults) are accounted for, greatly improving accuracy of immunoglobulin and biomarker estimates across all ages. It is a significant improvement over prior FT-IR diagnostic approaches which did not address age-dependent baseline shifts.

Thirdly, the invention provides a multi-output inference engine (708) capable of predicting multiple analytes and risk indicators from the single spectral input. After the spectrometer acquires a raw mid-IR spectrum, a preprocessing module (704) applies deterministic transformations – for example, polynomial baseline subtraction, Savitzky–Golay second-derivative filtering, and vector normalization – to produce a cleaned feature vector from the relevant spectral regions. This pre-processed spectral input (702) is then fed into the inference engine, which produces quantitative estimates of immunoglobulin levels (e.g., IgG and IgA concentrations) and key serum biomarkers such as an albumin level, a glucose-equivalent index, and total cholesterol, as well as qualitative or categorical outputs such as a sepsis risk indicator and a metabolic syndrome flag. All these results are generated simultaneously from the same spectral data. Decision thresholds for the classifications are adjusted to age-specific normal ranges or program-specific cut-offs. The outputs can be presented to the user in a single integrated report, for example with a traffic-light scheme (green/yellow/red flags) indicating normal vs. abnormal findings for each parameter (as illustrated in FIG. 5).

Importantly, the system includes a built-in quality control (QC) and outlier detection module (710) to ensure result reliability. The software computes statistical metrics such as Hotelling's $T^2$ and Q-residual for each new sample's spectral feature vector to measure how well it fits within the model's expected multivariate distribution. If a spectrum falls outside of the acceptable model domain (for instance, due to an unanticipated sample anomaly or instrument drift), the system suppresses or flags the outputs for that sample, preventing false results. Additionally, the instrument can enforce periodic two-level control checks using known reference sera: for example, every N samples the system may prompt insertion of low and normal concentration control specimens to verify that predictions remain within tolerance. A QC gating mechanism thus ensures that any significant deviation triggers an alert or recalibration before patient results are reported.

Another novel aspect of the invention is its federated self-calibration architecture for continuous learning across a network of deployed instruments. Each local device ("local bench" analyzer 302) can operate independently at the point of care, but it is also network-connected (securely, via encrypted channels) to a cloud-based analytics server (304). As illustrated in FIG. 4, the system implements a privacy-preserving federated learning loop: only aggregate summary statistics or model updates (and not raw patient spectra) are uploaded from the local bench (302) to the cloud. The cloud server aggregates data from multiple instruments and performs

periodic retraining of the master calibration models (e.g., using a federated averaging algorithm on accumulated gradient updates or calibration deltas). Updated model parameters (e.g., adjusted regression coefficients or new intercepts to correct drift) are then digitally signed and pushed back to each field instrument. This cloud-mediated self-calibration ensures that the analytical engine remains up-to-date and regionally calibrated as more data is collected, without compromising patient privacy. It addresses the challenge of regional biochemical diversity by gradually learning population-specific baselines: for example, the system's reference interval engine automatically recomputes age-stratified normal ranges (per guidelines like CLSI EP28-A3c) whenever a sufficiently large corpus (e.g., 10,000 samples) of apparently healthy spectra has accumulated. If the calculated normal range for any analyte shifts beyond a defined tolerance (for instance >0.25 standard deviations), the system can trigger a drift alarm for proactive review or recalibration. This closed-loop recalibration and reference interval governance module is a significant differentiator over prior art, which lacked automated, big-data-driven updating of diagnostic thresholds.

Furthermore, the invention includes an optional dual-channel decision logic (712) that fuses the spectroscopic analysis with external independent markers to improve diagnostic performance. In one embodiment, a Bayesian fusion module combines the probability outputs from the spectrometer's immunoglobulin prediction (which reflects humoral immune status) with results from cellular assays such as TREC/KREC counts or flow cytometry (which reflect T-cell and B-cell production) to compute a more robust overall PID risk score. For example, the system can ingest a TREC value (if available from newborn screening) and compute a posterior probability of "combined immunodeficiency" by statistically integrating both the spectral IgG result and the TREC result. This dual-channel approach mitigates the risk of false negatives that might occur if, say, maternal IgG in a newborn masks an antibody deficiency: the molecular TREC channel helps counterbalance the spectroscopic channel. The combined output can be used to triage infants who might otherwise pass a TREC-only screen but have emerging humoral immunodeficiency. This fusion of reagent-free spectroscopy with standard molecular tests in a single decision engine is unprecedented in the field of PID screening.

Overall, the present invention yields a rapid, reagent-free, and cost-effective screening tool. From a single tiny sample, it delivers multiple clinically actionable results (immunoglobulin levels, nutritional status indicator, lipid profile surrogate, etc.) within a minute, making it ideal for large-scale population health programs. The system is versatile, supporting both high-volume centralized laboratory use and deployment at remote clinics. In a central lab scenario, the device can be configured as a high-throughput analyzer with an automated conveyor feeding disposable sample tiles to the FT-IR spectrometer in rapid succession. In a field setting, the system can be embodied as a portable, ruggedized bench-top unit processing one sample at a time with minimal infrastructure. FIG. 1 provides a high-level block diagram of the system architecture, showing the integration of the sample interface, spectrometer, on-board computing, and network connectivity. The invention thus enables population-scale screening for humoral PIDs and NCDs that is both accessible and scalable, fulfilling a long-felt need in global health diagnostics.

# Brief Description of the Drawings

FIG. 1 is a block diagram of the overall system architecture, illustrating the flow from spectral data acquisition to preprocessing, age-based model selection, multi-task inference, and output generation (including quality control and flagging).

FIG. 2 is an exploded perspective view of an example disposable micro-volume ATR sampling tile and an automated conveyor mechanism. It shows the construction of the tile (602) and how a series of such tiles can be advanced to an ATR spectrometer interface for high-throughput operation, including a wiper/UV sterilization unit for cleaning between samples.

FIG. 3 is a schematic diagram of the data processing pipeline (software/AI flow), showing how a raw FT-IR spectrum is converted into results. The diagram highlights key stages such as spectral acquisition, signal preprocessing, feature extraction from specific wavelength windows, the age-adaptive model selection, an optional Bayesian dual-channel fusion module, and the decision logic that produces the final triage flags.

FIG. 4 is a flowchart of the cloud-based self-calibration loop (federated learning architecture). FIG. 4 depicts multiple local bench instruments (302) transmitting encrypted summary statistics to a central cloud server (304), which aggregates updates, retrains the master model, and returns updated calibration parameters to each instrument, thereby harmonizing their outputs over time without sharing raw patient data.

FIG. 5 is an example user interface or report output from the system, showing a one-page multi-analyte result. It includes quantitative readouts (e.g., IgG concentration, cholesterol level) with age-specific reference ranges (optionally presented as Z-scores) and simple traffic-light indicators (green/yellow/red) to flag normal or abnormal values, providing at-a-glance interpretation for clinical decision support.

FIG. 6 illustrates the micro-volume ATR sampling tile (602) in detail. It shows a top view and cross-section of the tile, including a hydrophobic boundary ring and hydrophilic center that self-limit the spread and thickness of a ~3 µL serum drop. This figure also indicates how the tile may be packaged in a sealed desiccant pouch (610) with a barcode (612) for identification and quality assurance.

FIG. 7 depicts an exemplary field sample preparation kit and workflow. It shows the use of a micro-capillary collection tube (604) to draw a finger-stick blood sample (approximately 5 µL), a portable battery-powered micro-centrifuge device (606) for rapidly separating serum from the micro-sample (yielding about 3 µL of serum in under 60 seconds), and the process of applying the serum to the ATR tile (602). FIG. 7 thus illustrates the practical steps by which an operator in a clinic would obtain a clear serum specimen on the tile for analysis.

FIG. 8A shows a workflow integration scenario in which the spectroscopic screening is piggy-backed onto an existing test. In this embodiment, the sample for the FT-IR system is obtained opportunistically during another routine procedure – for example, during an HIV rapid

diagnostic test (RDT) or a dried blood spot (DBS) collection for early infant diagnosis. The figure outlines how a small additional aliquot (less than 5 µL) can be taken from the same finger-stick without requiring a separate phlebotomy, and how this adds minimal time (under 60 seconds) to the existing process while not interfering with the primary test's reagents.

FIG. 8B illustrates an alternative high-throughput deployment of the system in a centralized laboratory setting. In this embodiment, a conveyor-belt or carousel mechanism continuously feeds ATR sample tiles into a bench-top FT-IR spectrometer for batch processing. Dual parallel belt feeds or other automation may be used to achieve throughput on the order of thousands of samples per day. FIG. 8B also shows that the core analytical engine remains the same in both the central lab configuration and the single-sample clinic configuration, underscoring the dual-deployment architecture controlled by one unified AI platform.

# Detailed Description of Preferred and Alternative Embodiments

## System Overview and Architecture

Referring to FIG. 1, the high-throughput spectroscopic diagnostic system comprises a combination of hardware and software components that work in concert to analyze a micro-volume blood sample and produce multiple diagnostic outputs. The hardware includes a local bench apparatus (302) which integrates a mid-infrared FT-IR spectrometer (310) and an on-board computing processor (312) within a portable or bench-top unit. The spectrometer (310) is configured to perform attenuated total reflectance (ATR) measurements in the mid-infrared range (for example, 4000–900 $cm^{-1}$) at a suitable resolution (e.g., ~4 $cm^{-1}$ nominal resolution). The local bench 302 further includes or interfaces with a sample presentation mechanism, as will be described below, to allow a small liquid sample to be placed in optical contact with the ATR crystal element of the spectrometer.

In the preferred embodiment for point-of-care use, the local bench (302) is a self-contained instrument approximately the size of a small desktop analyzer, requiring minimal external equipment. The on-board processor (312) can be an embedded single-board computer (for example, an ARM or x86 processor running a Linux OS) that handles spectral data processing, AI model inference, and user interface or data output tasks. The device may include a simple display or may connect to an external computing device for visualization of results. In alternative embodiments aimed at central laboratory deployment, the spectrometer and computing core can be integrated into an automated system with a conveyor or robotic sample handler for high-throughput operation (as illustrated in FIG. 8B). In either case, each local instrument (302) is optionally connected via a network to a cloud-based platform (304) for centralized data aggregation and model management (described later with reference to FIG. 4).

## Micro-Volume Sample Handling and ATR Tile

A critical hardware component of the invention is the disposable micro-volume ATR sampling tile (602), shown in FIG. 6. This tile serves as the sample carrier and optical interface between the biological specimen (serum or plasma) and the FT-IR spectrometer (310). The tile (602) is typically a small flat piece (for example, ~12 mm × 12 mm in size) with an ATR-active substrate, such as a diamond or germanium crystal, at its core. The top surface of the tile is engineered with a micro-patterned wetting area: a hydrophilic central region surrounded by a hydrophobic boundary. In use, a drop of the sample (approximately 3 µL of serum) is placed onto the tile's center. The surface chemistry causes the liquid to spread out evenly within the hydrophilic region but prevents it from overflowing beyond the defined area, thus creating a consistent thin film. This self-limited film has a controlled and reproducible thickness (target range about 2–5 µm) which is optimal for mid-IR absorption without spectral saturation. The uniformity of the sample layer on the ATR crystal leads to highly reproducible spectra across different tiles and operators, solving a common problem of pathlength variability in IR measurements of small volumes.

Each ATR tile (602) is intended for one-time use. To ensure stability and accuracy, the tiles are manufactured and packaged with care. They are stored dry in individual foil pouches (610) along with desiccant, preserving the surface from humidity and contamination. A barcode or RFID tag (612) on the pouch or tile can be used for tracking the lot, expiration, and calibration data of the tile. The packaging and tile design together guarantee minimal baseline drift (for example, ≤0.005 absorbance units over 30 days at elevated temperature) so that the system can maintain calibration over the shelf-life of the consumables. When a test is to be performed, a tile is removed from its pouch, the patient's serum sample is applied, and the tile is inserted into the instrument's ATR read position. After spectral acquisition, the used tile is ejected and safely discarded (for example, into a biohazard waste container inside the device).

In one preferred embodiment targeted at resource-limited clinics, the process of obtaining the serum sample for the tile is facilitated by a field-friendly kit (FIG. 7). A micro-capillary tube (604) is used to collect approximately 5 µL of capillary blood from a finger-stick. This micro sample is then immediately inserted into a small battery-powered micro-centrifuge (606). The micro-centrifuge (606) is a portable device – roughly akin to a "button" or coin-sized spinner – that can spin the capillary tube at high speed for a short duration (on the order of 30–60 seconds). Centrifugal force separates a few microliters of cell-free serum from the whole blood in that time. The design may include a wicking mechanism or simply allow the red cells to pack at one end of the capillary, leaving a clear supernatant. The operator then taps or wicks out about 3 µL of the serum and dispenses it onto the ATR tile (602). This field micro-separation protocol yields a usable serum sample rapidly without requiring a full laboratory centrifuge or large blood volume. The kit can be supplied as part of the system, including the capillary tubes (604), pre-packed ATR tiles (602) with pouches (610), the portable centrifuge (606), and step-by-step instructions for health workers. This enables the system to be used in clinics and screening camps with minimal training.

In an alternative embodiment for high-throughput settings, the sample handling is automated. FIG. 2 illustrates an example where multiple ATR tiles (602) are loaded onto a continuous conveyor belt or rotating carousel that passes under the spectrometer's ATR sensor

sequentially. Each tile receives a small serum drop (possibly via an automated pipettor) and is moved into position for measurement. After scanning, a built-in mechanism (such as a wiper arm applying 70% isopropanol and a burst of UV-C light at ~265 nm) sanitizes the ATR crystal or removes the used sample, and the spent tile is dropped into a waste bin. A fresh tile then moves into place for the next sample. Such an automated carousel could process on the order of one sample per second, translating to a theoretical throughput of up to ~3600 samples/hour; a more practical implementation might achieve ≥1000 samples/hour (taking into account handling overhead). This high-speed, conveyor-based ATR sampling subsystem has not been described in prior art for biological fluids, and it allows the invention to scale for large population screening programs or central lab workflows.

## Spectral Acquisition and Preprocessing Pipeline

Once the sample is in place on the ATR crystal, the FT-IR spectrometer (310) collects a spectrum. In a typical implementation, the instrument first records a background spectrum (e.g., averaging 32 scans with no sample or with a reference material) to serve as a baseline. Then it records the sample spectrum (for instance, averaging 10 scans of the sample droplet at 4 $cm^{-1}$ resolution across the mid-IR range 3900–900 $cm^{-1}$). The raw output is an interferogram transformed into an absorbance spectrum, which contains peaks corresponding to various molecular bonds in the sample.

The system's software then executes a deterministic preprocessing pipeline (704) on the raw spectral data. This step is important for normalization and feature extraction, especially given the complex composition of human serum. In the preferred embodiment, the preprocessing (illustrated in the flow of FIG. 3) involves the following sub-steps:

- Band-limiting the spectrum: Only specific regions of the spectrum that carry relevant biochemical information are retained for analysis. For example, the system may restrict the analysis to a "protein window" around 1750–1300 $cm^{-1}$ (covering the amide I and II bands of proteins and the $CH_2$ wag region related to lipids) and a "carbohydrate/small-molecule window" around 1150–950 $cm^{-1}$ (which contains signatures from sugar and glycoprotein vibrations). By focusing on these bands, the influence of irrelevant regions (such as the strongly absorbing water bands or empty regions with no signal) is minimized. FIG. 3 conceptually shows these spectral windows being selected from the full spectrum.

- Baseline correction: A polynomial baseline is subtracted from the spectrum to correct for any sloping background (due to, e.g., scattering or instrument drift). This step (sometimes referred to as detrending) ensures that the subsequent analysis focuses on the true absorption peaks rather than on baseline offset.

- Derivative filtering: A Savitzky–Golay smoothing filter is applied to compute the second derivative of the spectral signal (within the chosen windows). Using a second-derivative (with, for instance, a 9-point window and 2nd-order polynomial fit) enhances resolution by sharpening spectral features and removing constant offsets. It also tends to normalize

the spectrum in terms of peak width, making the analysis less sensitive to broad baseline changes. This is a common chemometric technique to handle overlapping peaks in IR spectra.

- Vector normalization: The processed spectral values are then $\ell_2$-normalized (unit-length normalization) within each retained band (or across the concatenated bands) to account for pathlength or concentration differences. Essentially, the intensity vector of the spectrum is scaled so that its magnitude is consistent (e.g., unity), which helps the model focus on the shape of the spectral signature rather than absolute intensity. This also mitigates any remaining variations in sample thickness on the tile.

After these steps, the result is a cleaned and standardized feature vector (702) representing the essential spectral information of the sample. In one example embodiment, the final feature vector may have 451 data points (after concatenating multiple spectral regions post-derivative and down-sampling), though the exact length can vary based on how the spectral windows are defined and processed. This feature vector (702) is then passed into the machine learning inference stage.

## Age-Adaptive Chemometric Engine and Decision Logic

One of the core innovative elements of the system is the age-adaptive chemometric engine (706) which interprets the preprocessed spectral data to produce analytic results. Because normal levels of immunoglobulins and other biomarkers change with age (especially from infancy through adulthood), the engine is designed to adapt its model based on the age of the subject.

Upon receiving the spectral feature vector, the system also takes as input the subject's age (which could be entered by the user or obtained from patient records). An age-band selector within module 706 uses this age to route the sample to the appropriate calibration model. In an exemplary implementation, the system might maintain four calibration models for different age ranges: e.g., 0–6 months, 6–24 months, 2–5 years, and >5 years (including adults). These models can be, for instance, partial least squares (PLS) regression models or neural network models trained on reference data for the respective age group. The age-band selector computes an index (for example, by dividing age in months by a predefined band width and flooring it) to select the correct model coefficients (β-vector or trained network weights) corresponding to that age group.

The multi-output inference engine (708) then uses the selected model to transform the spectral features (702) into multiple predictions. If implemented as a PLS regression, this could be a PLS2 model that yields multiple dependent variables; if implemented as a machine learning model, it could be a multi-task neural network with several output nodes. In the preferred embodiment, the outputs include at least:

- Quantitative concentration estimates for immunoglobulin G (IgG) and immunoglobulin A (IgA) (measured in, say, mg/dL). These serve as the humoral immune markers to detect immunoglobulin deficiencies.

- Quantitative estimates for metabolic markers such as albumin, glucose-equivalent (an inferred glucose level or a correlated glycemic index), and total cholesterol. These represent nutritional and NCD-related analytes derivable from the serum's IR signature.

- Binary or ordinal classification outputs for conditions like sepsis risk and metabolic syndrome presence. These are derived from combinations of spectral features (for example, a high sepsis risk might be flagged if certain inflammation-related spectral features are extreme).

- Optionally, other surrogate indices such as a renal function indicator (which could be a composite index reflecting urea/creatinine levels gleaned from certain spectral bands) can be included.

The inference engine uses the spectral patterns in the chosen IR windows (e.g., the amide I/II region for protein content, the 1150–950 $cm^{-1}$ region for carbohydrate content, and a lipid-associated absorbance around 1740 $cm^{-1}$ for cholesterol/triglycerides) to predict these outputs. In FIG. 3, this multi-task prediction step is depicted where the single feature vector fans out to multiple result values.

After obtaining the raw predictions, the system applies its decision logic and QC gating. The quality control module (710) evaluates the Hotelling $T^2$ statistic (a multivariate measure of how far the sample's feature vector is from the training set center) and the Q-residual (how much of the sample's variance is unexplained by the model). If either metric exceeds predefined thresholds, it indicates the sample is an outlier or of poor quality. In such a case, the logic may invalidate the results for that sample – for example, by not displaying numerical values and instead showing an error or "QC fail" flag prompting a re-test. If the metrics are within range, the results are considered valid.

The decision logic also compares the numeric predictions to reference ranges or cut-offs relevant for the subject's age. The system includes an internal reference interval database stratified by age (and potentially by region or population, as updated by the cloud engine described later). For each analyte or risk output, the system can determines if it falls in the normal range, moderately abnormal, or severely abnormal for that age. It then assigns a color-coded flag or indicator accordingly (for instance, green for within normal limits, yellow for mildly out-of-range, red for critically out-of-range). These flags, along with the numeric values and unit labels, are compiled into the output report.

The final compiled results can be delivered to the user through various means. In a stand-alone device scenario, the local bench (302) may simply display the results on a screen or print them. In more connected implementations, the results can also be packaged into a standard digital

form (for example, a single HL7/FHIR compliant message (as per one embodiment) to integrate with electronic medical record systems). A report generator (720) produces the output for display or transmission. FIG. 5 exemplifies how the outputs might be presented, with each analyte accompanied by a gauge or colored icon indicating the interpretation. By consolidating all the information into one report, the system simplifies the clinical decision process – a healthcare provider can immediately see if the patient likely has a normal immune profile, or if, for example, low IgG (red flag) and low albumin (red flag) are present, suggesting further investigation into immunodeficiency and malnutrition is warranted.

## Dual-Channel Fusion with External Markers

In an enhanced embodiment, the system optionally incorporates a dual-channel diagnostic fusion (712), which combines its spectroscopic analysis with external test results to improve overall screening accuracy. This is particularly useful for comprehensive PID screening. As noted, newborn screening for severe combined immunodeficiency (SCID) uses molecular tests like TREC (T-cell receptor excision circle) counts to catch T-cell defects. However, solely relying on TREC misses B-cell or antibody deficiencies; conversely, an FT-IR spectroscopic test focusing on Ig levels could miss certain cellular immunodeficiencies or be confounded in newborns by maternal IgG.

To address this, the invention can ingest data from external assays — for example, TREC and KREC (Kappa-deleting recombination excision circle) counts obtained via PCR, or flow cytometry results for T-cell counts — and integrate them with the spectroscopic findings. In practice, a Bayesian network or Bayesian inference engine (712) takes as inputs the probabilities or risk scores from the spectrometer (e.g., probability of PID based on low Ig levels) and the probabilities from the other tests (e.g., probability of PID based on low TREC, etc.). FIG. 8A conceptually illustrates this approach. The engine applies Bayes' rule or logistic regression on these inputs, using pre-learned weighting, to compute a combined posterior probability of a primary immunodeficiency. By using a probabilistic fusion, the system accounts for the independent evidence provided by each channel.

For instance, if the spectroscopic channel alone yields a borderline-low IgG (slightly elevated risk) but the TREC channel is very low (high risk), the combined posterior risk will be higher than either alone, prompting referral despite the borderline IgG. On the other hand, if one channel is likely a false alarm (e.g., low TREC but normal Ig and other markers), the fusion can reduce false positives by requiring concordance or by weighting each channel appropriately. The outcome of this dual-channel analysis is presented as an overall PID risk score (0 to 1) and possibly categorized into a traffic-light category (e.g., green = low risk, red = high risk requiring follow-up). Incorporating such dual-channel logic is novel in that no prior diagnostic platform fuses spectroscopic data with molecular test results in one system. This provides a more holistic assessment of an individual's immune status and can be especially valuable in newborn screening programs or immunology clinics, as it leverages the strengths of multiple methodologies.

## Federated Learning and Cloud-Based Calibration Updates

With reference to FIG. 4, the invention employs a cloud-connected data system (304) to continually improve and maintain the accuracy of its diagnostic models across a distributed network of devices. Each local bench unit (302) operates independently for real-time analysis, but it also participates in a collaborative model updating scheme. After each test (or periodically in batches), the device packages summary information about the run – for example, the sample's feature vector projections, the initial predictions, and any available ground-truth confirmations (if later provided by laboratory tests) – in a secure, anonymized format. Only high-level summary statistics are transmitted, such as aggregated residuals or gradient snippets, ensuring that no personally identifiable raw data leaves the device.

The cloud server (304) aggregates data from numerous deployed instruments (302 across different regions or hospitals). On a scheduled basis (e.g., nightly), it performs a federated learning update. In one implementation, it uses a federated averaging (FedAvg) algorithm: it computes incremental model adjustments (deltas) from the summary data of each site, averages these updates, and applies them to the master model. The cloud may retrain parts of the calibration – for example, adjusting the regression intercepts and slopes to align with new reference lab values that have been reported, or expanding the model's training set with spectra from diverse populations to improve its robustness.

Once a new set of model parameters is obtained, the server digitally signs the update (to ensure authenticity and integrity) and distributes it back to all the bench units (302). The benches then update their local model library (706) with the new coefficients. This way, even if, say, a particular clinic starts encountering patients with an atypical biochemical profile (perhaps due to regional diet or genetics), the model can adapt globally to that data over time. Additionally, the cloud platform recalculates the reference intervals for each analyte using the growing dataset of presumably healthy cases (implementing an indirect reference interval method as per CLSI EP28-A3c). For example, after every 5,000 new normal samples collected, the cloud module may recompute the 2.5th and 97.5th percentile for IgG for each age band. These updated normal ranges are then sent to devices so that their flagging criteria remain accurate.

This cloud-based reference interval governance and drift detection is another innovative aspect. If the cloud notices a systematic drift (e.g., all devices in a region gradually reporting lower albumin, indicating a calibration bias or reagent lot issue), it can alert the operators or automatically adjust the calibration curve to correct the intercept. All such updates and data exchanges are conducted under strict security: communications can use TLS encryption and devices can contain hardware security modules to verify the authenticity of model updates (for compliance with regulations such as FDA 21 CFR Part 11 for electronic records). The net effect is a self-learning network of analyzers that gets "smarter" and more accurate with each day of use, without manual re-calibration by technicians.

## Anomaly Detection and Regional Personalization

Human populations have variability in diet, genetics, and environment that can lead to region-specific spectral patterns. To ensure the system works reliably as it is deployed broadly, the invention includes an anomaly detection and regional personalization layer. This can be implemented via unsupervised machine learning techniques (for example, an autoencoder or one-class SVM, or by clustering methods like t-SNE/UMAP on the spectral dataset). The idea is to identify when a sample's spectral signature does not resemble any data seen in training or in the accumulating cloud database – such a sample might indicate a rare condition or an error. The anomaly detector runs either on the device or in the cloud and assigns a score to each spectrum for how "out-of-distribution" it is relative to the learned model. If this score is high, the system flags the sample as a potential regional outlier. It might then, for instance, suggest a reflex confirmatory test (i.e., advise sending the sample to a lab for confirmatory testing because the spectral reading is unusual).

Additionally, the model library (706) can be regionally tailored: the cloud might maintain slightly different model parameter sets for different clusters of population data if it finds significant differences. For example, in some regions with prevalent malaria or other conditions that alter serum composition, a separate calibration curve might perform better. The federated learning framework can accommodate this by clustering updates or weighting them differently. In one embodiment, the cloud gives higher weight to certain loss functions based on context (e.g., in pediatric cohorts, prioritize PID-related prediction accuracy, whereas in adult cohorts, prioritize metabolic syndrome predictions). This dynamic adjustment ensures the personalization of the screening tool to the local population without requiring a complete re-development per region.

## Deployment Modes and Use Cases

The versatility of the invention allows it to be deployed in multiple modes. Two primary deployment embodiments have been contemplated:

1. Central Lab High-Throughput Analyzer (FIG. 8B): In this mode, the system is installed in a laboratory or screening center with access to mains power and possibly robotics. It operates with a dual-belt or automated feed system, processing thousands of samples per day. This configuration might be used by national reference labs or large hospital labs to run population-scale screening (for example, adding a newborn immunoglobulin screening panel to existing newborn screening programs). The device in this setting could be paired with on-site confirmatory testing equipment (like a nephelometer or flow cytometer) for any positive cases, streamlining the workflow.

2. District Clinic or Point-of-Care Bench (FIG. 7): In this mode, a single-sample, ruggedized bench-top unit (302) is deployed at a smaller clinic, health post, or mobile screening camp. It runs on standard AC power or a backup battery/UPS if power is unreliable. It communicates via cellular data or satellite (e.g., 4G or Starlink) to send data to the cloud if available, but can also function offline. This unit uses the micro-capillary and mini-centrifuge kit for sample prep, requiring only minimal training for health workers. It provides immediate results that can guide decisions such as whether to refer a patient for further immunological evaluation or to initiate nutritional interventions. Because it

uses no wet reagents (just the dry ATR tile), it is robust and cost-effective in low-resource settings and does not generate biohazardous liquid waste.

A key advantage of the invention is that both deployment types use the same core AI engine and methods. A screening program could thus implement central labs in urban areas and bench units in rural clinics, all feeding data into the same cloud system (304). This ensures consistency in results and allows improvements learned in one context to benefit all others via federated updates.

Another use case illustrated in FIG. 8A is leveraging existing healthcare touchpoints. For example, during routine immunization or maternal-child health visits, a finger-prick blood sample is often collected for various tests (such as HIV early infant diagnosis via DBS, or malaria rapid tests, etc.). The present system can be integrated into those workflows by utilizing the same finger-prick sample to perform the FT-IR analysis. This "piggy-back" approach means the screening for PIDs and NCD markers can be added without extra phlebotomy and with negligible additional time. The claimable benefit here is that the invention can be implemented opportunistically during standard care procedures (like an HIV RDT), making population screening highly efficient.

## Example Operation

In practice, a typical operation of the system for a patient sample would proceed as follows: A health worker collects a drop of blood via finger-stick and uses the kit's micro-capillary (604) and centrifuge (606) to obtain serum, which is then dropped onto an ATR tile (602). The tile is inserted into the spectrometer slot of the bench device (302). The device, perhaps via an integrated sensor, detects the presence of the tile and initiates a measurement. The FT-IR spectrometer (310) scans the sample and produces a raw mid-IR spectrum. The on-board processor (312) then executes the preprocessing steps (baseline removal, derivative, normalization) in module 704, yielding the feature vector (702). The subject's age (e.g., 9 months old) is input, either manually or from patient data, and the age-adaptive engine (706) selects the appropriate infant calibration model. The model is applied in the inference engine (708) to predict, for instance, IgG = 250 mg/dL, IgA = 20 mg/dL, albumin = 3.8 g/dL, glucose-equiv = 100 mg/dL, total cholesterol = 160 mg/dL. It also outputs classification scores, e.g., sepsis risk = "low" (green), metabolic syndrome = "N/A" (since an infant, this might not apply).

The QC module (710) calculates that this sample's $T^2$ and Q values are within normal range, so it does not trigger a QC fail. The reference comparison engine then checks the age-specific normal ranges: for a 9-month-old, suppose normal IgG is 300–700 mg/dL; the measured 250 mg/dL is below the 2.5th percentile for age, so that result is flagged (red) as low IgG (potential immunodeficiency). IgA might also be low but IgA in infants is naturally low, so its flag might be green if within expected infant range. Albumin 3.8 g/dL is within normal for age (green), glucose 100 mg/dL is normal (green), cholesterol 160 is normal (green). The device displays a report indicating low IgG with a red highlight and suggests "PID Triage: Positive – low IgG" while other

NCD markers are normal. If available, the worker might input a TREC test result (for example, from the national newborn screening card). If the TREC was normal, the dual-channel fusion (712) might still indicate a moderate combined PID risk due to the IgG result, but less severe than if both were abnormal. The final triage recommendation might be to refer the infant for further immunology evaluation (given the humoral deficiency signal).

After this, the device (302) optionally logs a summary of the case (e.g., the fact that IgG was low and maybe a hash of the spectrum) and sends anonymized data to the cloud (304) when network connectivity is available. Over that night, the cloud might include this case (especially since it was flagged) in the next model refinement, thereby gradually improving the system's performance.

Through this comprehensive approach, the invention provides a powerful diagnostic tool that can drastically reduce the time and resources needed to screen for a range of conditions. It brings together hardware innovation (micro-sampling and high-throughput ATR), software innovation (age-adaptive multi-task AI with QC gating), and system innovation (federated learning and workflow integration) into a single platform. The result is an apparatus and method that enable population-scale quantitative immunoglobulin triage alongside screening for common chronic diseases, in a manner that is far more accessible and rapid than conventional lab testing.

# Reference Numeral Index

- 302 – Local bench apparatus (point-of-care FT-IR spectrometer device, including internal spectrometer and processor)
- 304 – Cloud server or remote data hub (for federated learning and central model updates)
- 310 – FT-IR spectrometer unit (mid-infrared attenuated total reflectance spectrometer within the bench device)
- 312 – On-board processor (computing unit in the local bench for data analysis and control)
- 602 – Micro-volume ATR sampling tile (disposable sample carrier with hydrophilic/hydrophobic pattern for ~3 µL serum)
- 604 – Micro-capillary tube (capillary blood collection tube for ~5 µL finger-stick sample)
- 606 – Portable micro-centrifuge (battery-powered device for rapid field separation of serum from micro blood samples)

- 610 – Desiccated foil pouch (protective packaging for the ATR tile to maintain stability and cleanliness)
- 612 – Barcode or identifier (label on the tile or pouch for tracking and calibration reference)
- 702 – Spectral input vector (preprocessed FT-IR spectral data used as input to the analysis engine)
- 704 – Preprocessing module (software component performing baseline removal, derivative filtering, normalization on the spectrum)
- 706 – Age-adaptive model selection engine (module that chooses the appropriate calibration model based on subject age)
- 708 – Multi-output inference engine (AI/chemometric model producing multiple analyte predictions and risk classifications from the spectral input)
- 710 – Quality control and outlier detection module (monitors spectral consistency with model, computes $T^2$/Q stats, and gates/suppresses unreliable results)
- 712 – Dual-channel fusion module (Bayesian or probabilistic engine fusing spectroscopic outputs with external test data like TREC/KREC for combined diagnosis)
- 720 – Report generator / output interface (component that compiles results and flags into a user-friendly report or data message)

# Figures