Natural Resources for Human Health

eISSN: 2583-1194

Original Research

Received 20 March 2026
Revised 15 April 2026
Accepted 14 May 2026

Natural Resources for Human Health 2026, Vol. 6 (2s): 82–91

KEYWORDS:
Blockchain; Federated Learning; Public Health Surveillance; Pandemic Intelligence; Privacy Preservation; Explainable Artificial Intelligence.

A Secure Blockchain-Enabled Federated Learning Framework for Privacy-Preserving Public Health Surveillance and Pandemic Intelligence

[1]Dr. Keerthika Velusamy, [2]Dr. K. Sundareswari, 3S. Syedzagiriya, [4]Mr.M.Mohideen Batcha, [5]Dr. Sankar K, [6]Dr. Sumit Kushwaha, [7]R. Naveenkumar, [8]Ramesh Mylapalli,

[1]Associate professor & Head., Department of Cybersecurity, KPR College of Arts Science and Research, Coimbatore, keerthiramesh83@gmail.com Orcid ID: 0000-0002-8817-9978
[2]Professor in Mathematics, Al-Ameen Engineering College (Autonomous), Erode, Tamilnadu, India. sundarimaths@gmail.com 0000-0002-9642-5657.
[3]Assistant Professor, Department of ECE, Al-Ameen Engineering college (Autonomous), Erode, Tamilnadu, India Email id: zagiriya@gmail.com Orchid id: 0009-0002-1976-5065
[4]Assistant Professor / Mechanical Engineering, Al-Ameen Engineering College (Autonomous), Erode. Mail Id: mmb4mech@gmail.com Orcid: 0009-0008-8099-4717
[5]Professor, Community Medicine, Meenakshi Medical College Hospital & Research Institute, Meenakshi Academy of Higher Education and Research, Enathur, Kanchipuram, Tamil Nadu 631552
[6]Associate Professor, Department of Directorate of Research, Innovation, and Consultancy (DRIC), Chandigarh University Uttar Pradesh, Unnao - 209859, Uttar Pradesh, India. Email: sumit.kushwaha1@gmail.com ORCID: 0000-0002-3830-1736
[7]Dept of CSE, School of Engineering and Technology, CGC University Mohali-140307, Punjab India.Email: drnk1983@gmail.com 0000-0001-9033-9400
[8]Assistant Professor, Department of Computer Science & Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur District, Andhra Pradesh, India mylapalli2@gmail.com https://orcid.org/0009-0007-6919-4156

ABSTRACT: Public health monitoring demands the quick enable of multi-institutional intelligence – but not patient-sensitive data. The following paper introduces a federated learning system based on the blockchain for pandemic surveillance in 10 healthcare institutions while preserving user privacy. The framework is comprised of two-layer LSTMs based on 35 demographic, clinical, physiological, exposure, vaccination, and community-level variables with an observation window of 14 days and a prediction horizon of 7 days. Over 100 communication rounds, there are 5 local epochs in which institutions do not transfer patient level records. Model updates are protected by a trust threshold of 0.60 and differential privacy budgets ranging between ε=0.5 and ε=8, secure aggregation, hashing with SHA-256, and ECDSA-256 signatures. ROC-AUC of $\geq$ 0.85, sensitivity of $\geq$ 85%, poisoning-detection of $\geq$ 90%, invalid-update rejection of $\geq$ 95%, subgroup performance differences of $\leq$ 10 percentage points, and blockchain latency of $\leq$ 2 seconds are all among the evaluation targets. The framework generates individual risk probabilities, 7-day forecasts of outbreaks, alerts for outbreaks and clusters, explanations using SHAP and estimates of healthcare resources, setting measurable needs for future validation.

1. INTRODUCTION

Efficient and timely pandemic surveillance is key for the ability to detect emerging outbreaks, monitor spread, guide testing and limited health-care resources. Hospitals, laboratories, schools and public health agencies routinely provide data that are stored in a centralized database. While centralization makes analysis easier, it also makes for targets that can be attacked, accessed by unauthorized individuals, tapped into for data mining without authorisation or even reused in other surveilling ways than the one the data was originally used for. Young people have particular concerns about these findings because their medical history can contain indicators such as symptoms, vaccination records, comorbidities, school attendance, household exposure and socioeconomic information, along with indicators on location. These data must be protected to maintain confidentiality, accountability for the institution and trust among the public.

However, with federated learning, institutions can jointly train a prediction model without sharing patient-level records. However, traditional FL is susceptible to gradient leakage, member-inferencing attacks, model poisoning, unreliable member participation, and poisoning by a corrupted aggregation server. It also offers only partially mechanisms for asserting model provenance and maintaining audit histories of changes that give a trace of a model's changes, and ensure that these changes are not tampered with.

To overcome these drawbacks, this paper presents a secure federated learning system based on the blockchain technology, which combines the advantages of differential privacy, secure aggregation, trusted evaluation of updates and explainable temporal prediction. Differential privacy reduces the leakage of information from local updates, secure aggregation obscures the individual institutional parameters, and permissioned blockchain provides identity, signatures, update hashes and model versions verification. The goal of the framework is to achieve the following benefits while keeping sensitive records decentralized: 1) Risk estimation of new infections over seven days, 2) regional outbreak forecasting, 3) cluster detection, and (4) planning for healthcare resources. The proposed architecture addresses prediction, privacy and blockchain problem in an integrated way rather than having the three considered separately, as in other approaches. It has made contributions such as a multi-institutional learning architecture, a blockchain-based validation workflow, a defence mechanism against malicious updates through trust, and a multidimensional evaluation protocol based on accuracy, calibration, privacy, security, fairness, communication efficiency, scalability and operational reliability requirements.

2. RELATED WORK

Classification, temporal forecasting, and anomaly-detection models are an integral part of machine learning and have proven useful for pandemic surveillance to detect infection risks, forecast transmission trends, recognize emerging clusters, and predict the need for sources of healthcare. Typically, models are trained with clinical and epidemiological data collected at the central level. While these methods can enable quick analysis, these rely on a centralized database, which raises privacy, governance and cyber security issues, especially if records include data on adolescents, household exposure, or geographical behaviour, and school attendance [12, 13].

In the medical field, federated learning has thus been explored as an approach to learn models at the shared level without having to share patient-level information among hospitals [6, 10, 12, 13]. Some studies have utilized federated learning for the diagnosis of disease, medical imaging, electronic health records, and forecasting outbreaks [3, 4, 12]. But storing information on your local device is not 100% private. Gradient and model parameters can leak membership data, or allow partial reconstruction of the training records [8]. Federated systems are also vulnerable to poisoning [8] and backdoor [11], inference and model-replacement attacks from compromised participants.

Doctors and patients can trust the integrity and accountability of a blockchain-based healthcare architecture using distributed identity management, immutable audit trails, smart contracts and cryptographic verification [1, 7]. However, a direct transfer of health-related sensitive data to the ledger increases security and data integrity concerns. Both differential privacy [5] to bound the influence of individual records, and secure aggregation [2] to avoid allowing the coordinator to view individual institutional updates, exist in the literature. These protections can diminish the ability to predictively use the data, add to communication overhead, and not recognize semantically valid but malicious contributions [6], [8], [11].

Explainable AI (XAI) has been applied to translate key symptoms, clinical characteristics and community signals that lead to public-health predictions [9]. However, descriptions of the statistical model behaviour usually do not give any explanation on disease mechanisms. Current frameworks mainly focus on one aspect of prediction, federated privacy, the integrity of

blockchain, or explainability, and address other aspects with limited evaluation [6, 7, 12]. The research gap is therefore a combined and easy-to-trace surveillance system able to protect and manage spread records, ensure participant verification, detect malicious updates, facilitate interpretable forecasting and assess operational trade-offs.

3. PROPOSED METHODOLOGY

3.1 Overall Framework

The proposed structure offers possibilities to create a model for a pandemic surveillance system collaboratively among hospitals, laboratories and public health bodies without transfer of patient level data. It includes six steps: participating institutions, local preprocessing, federated LSTM training, privacy protection, privacy protected federated aggregation, and pandemic-intelligence generation, as illustrated in Figure 1.

Every institution has local records, provides missing value filling, normalizes 35 input variables, and produces sequences of 14-days of observations. During each communication round a two-layer LSTM is trained locally for five epochs. Transmission of only protected model updates.

The differential privacy restricts leakage of information, and the secure aggregation hides individual institutional update. A permissioned blockchain validates institutional credentials, digital signatures, SHA-256 hash, model versions and submission time. Any updates with a trust score of 0.60 or higher are allowed to be aggregated worldwide.

The created global model is spread back to the participating institutions for further training sessions. Once converged, it provides seven-day infection-risk assessments, regional outbreak predictions, cluster notifications, resource-needs projections and SHAP explanations. Figure 1 as a result provides a seamless and traceable process for privacy protection, collaborative learning in a secure environment, and pandemic intelligence.

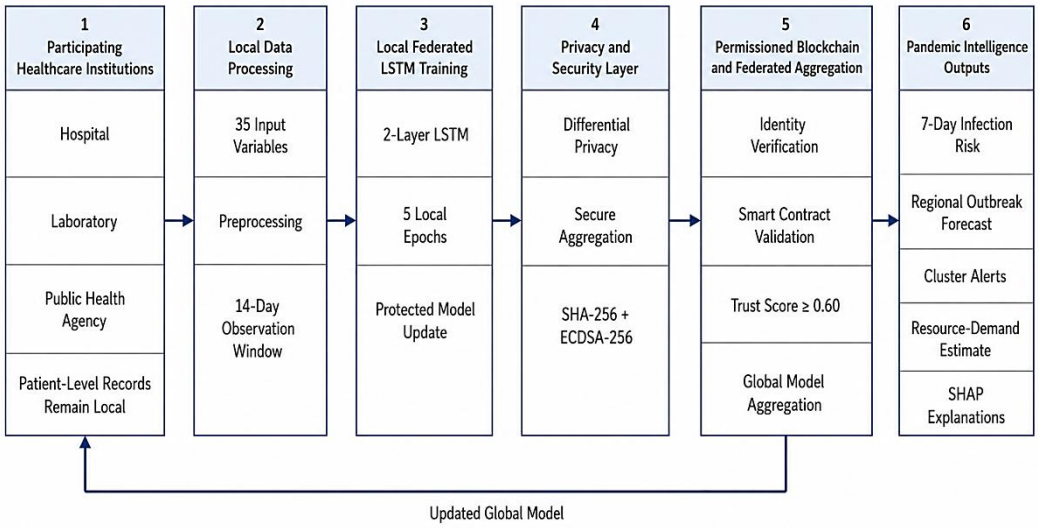

Figure 1. Proposed blockchain-enabled federated surveillance framework.

3.2 Inputs, Preprocessing and Temporal Prediction

The proposed model handles 35 variables of demographic characteristics, symptoms, physiological measurements, comorbidities, exposure history, vaccination status, community conditions and temporal information. These variables provide clinical information at individual level as well as epidemiological context at population level, as summarised in Table 1.

Categorical data values indicating missing data are filled in with the category "unknown", and continuous data values are filled in with the median value calculated by each participating institution. When there is a potential predictive information in the absence of a measurement, then the missingness indicator is kept. Continuous features are normalized based on their local training set mean and standard deviation and categorical variables are encoded using one-hot encoding.

Observations are preprocessed and presented in 14-day long temporal windows and passed through a two-layer LSTM network with 128 and 64 units of hidden layers. The model predicts a defined infection-related outcome in the next 7 days in the adolescent (i) when the outcome is not certain. The temporal risk prediction is seen as the following equation:

$$\hat{y}_{i,t+7} = \sigma\left(W_o h_{i,t}^{(2)} + b_o\right), \quad (1)$$

where $h_{i,t}^{(2)}$ is the representation generated by the second LSTM layer, W_o and b_o are the output-layer weight and bias, and $\sigma(\cdot)$ is the sigmoid function. The resulting value $\hat{y}_{i,t+7} \in [0,1]$represents the seven-day infection-risk probability. Patient-level outputs include risk probability, risk category, and contributing factors, whereas population-level outputs include outbreak forecasts, cluster alerts, incidence trends, and resource-demand estimates.

Table 1. Input and output variables of the temporal surveillance model

Category	Included variables or outputs	Number
Demographic inputs	Age, sex, residential setting, socioeconomic category	4
Symptom inputs	Fever, cough, fatigue, sore throat, dyspnoea, headache, anosmia, gastrointestinal symptoms	8
Physiological inputs	Temperature, heart rate, respiratory rate, oxygen saturation	4
Comorbidity inputs	Asthma, obesity, diabetes, cardiovascular condition, immunocompromised status	5
Exposure inputs	Household contact, school exposure, community exposure, recent travel	4
Vaccination inputs	Dose count, vaccine category, time since last dose	3
Community inputs	Regional incidence, positivity rate, absenteeism, hospital occupancy	4
Temporal inputs	Surveillance day, epidemiological week, symptom duration	3
Total inputs	**Combined surveillance variables**	**35**
Patient-level outputs	Seven-day risk probability, risk category, contributing factors	3
Population-level outputs	Outbreak forecast, cluster alert, incidence trend, resource demand	4

3.3 Privacy-Preserving Federated Learning

In every communication round, the coordinating server shares the current global model to certain health care institutions. The model is trained locally in each institute for 5 epochs with its own records. The patient-level information is kept in the originating institution, with only protected model parameters sent to be aggregated.

Each local update is clipped to a norm of 1.0 and Gaussian noise is added that is calibrated according to differential privacy. The privacy-utility trade-off is evaluated on privacy budgets of between $\varepsilon = 0.5$ and $\varepsilon = 8$ with $\delta = 10^{-5}$. Cryptographic aggregation hides the local parameters in such a way that the coordinating server can only read their sum.

The proposed approach assigns a weight to each of the accepted institutional models based on the number of samples collected in that location (local sample size) and confirmed trust score. Institutions with a higher trust score have a higher influence (as calculated in Equation (2)), while updates with a lower score have a lower influence.

$$\theta^{(r+1)} = \sum_{k=1}^{K} \frac{n_k T_k}{\sum_{j=1}^{K} n_j T_j} \theta_k^{(r+1)}, \quad (2)$$

where $\theta^{(r+1)}$ is the updated global model, $\theta_k^{(r+1)}$ is the local model submitted by institution k, n_k is its number of eligible training records, $T_k \in [0,1]$is its validated trust score, K is the number of accepted institutions, and (r) denotes the communication round. Only updates that have been verified as being free from privacy, identity, integrity and trust issues are included. The overall federated training process is summarized in Figure 2 and the key model and privacy parameters are shown in the Table 2.

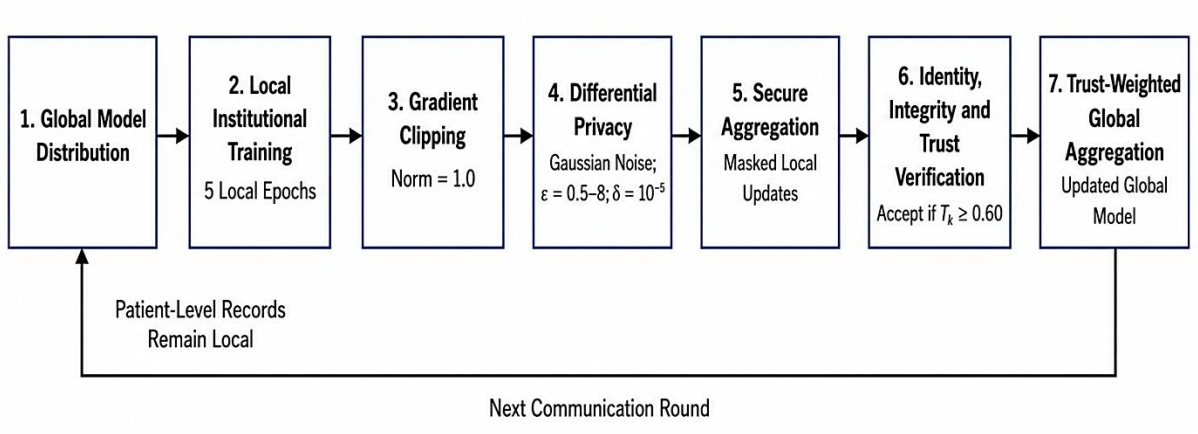

Figure 2. Privacy-preserving federated training process.

Table 2. Model and privacy parameters

Component	Parameter	Proposed value
Input configuration	Input variables	35
Temporal model	Observation window	14 days
Temporal model	Prediction horizon	7 days
Temporal model	LSTM layers	2
Temporal model	Hidden units	128 and 64
Temporal model	Dense-layer units	32
Temporal model	Dropout rate	0.30
Local training	Local epochs per round	5
Local training	Batch size	64
Local training	Learning rate	0.001
Local training	Optimizer	Adam
Federated training	Maximum communication rounds	100
Federated training	Minimum institutional participation	70%
Differential privacy	Privacy budgets	0.5, 1, 2, 4 and 8
Differential privacy	Delta	10^{-5}
Differential privacy	Gradient-clipping norm	1.0
Secure aggregation	Minimum participation threshold	70%
Trust validation	Trust-score range	0–1
Trust validation	Minimum accepted trust score	0.60

3.4 Blockchain and Trust-Based Security

The framework proposed is a permissioned blockchain, which means that only registered healthcare institutions, laboratories and public health agencies can be part of it. The network membership authority (NMA) will provide each institution with a verified digital identity and a pair of key pairs: public key ECDSA-256 and private key ECDSA-256. The institution first

produces a hash of the local model using the SHA-256 algorithm, binds its own digital signature to it, values the model version, adds the number of the communication round, and adds a time stamp before submitting a local model update.

The institutional identity, digital signature, updates hash, model version, submission time and duplication status are automatically checked by smart contracts. Patient information and raw parameters of the models are not stored on the blockchain. Only various update hashes, institutional IDs, trust scores, validation results, and model-version information are stored on the ledger.

Malicious-update detection relies on the similarity of the models submitted, the performance of the models in the test set, the historical trustworthiness of the institutions and the internal integrity of the cryptographic signatures. The institutional trust score is given by Equation (3):

$$T_k = 0.30S_k + 0.35V_k + 0.20H_k + 0.15I_k, \quad (3)$$

where T_k represents the trust score of institution k, S_k is the similarity between its local update and the accepted global-update direction, V_k is its validation performance, $H_k$represents historical reliability, and I_k indicates successful identity, signature, and integrity verification. All components are normalized between 0 and 1.

As specified in Equation (3), an update is accepted for global aggregation when$T_k \geq 0.60$. Any updates below are withheld and submitted for further inspection. The blockchain validation workflow from institutional authentication, model hashing to smart-contract verification, trust assessment decision and final acceptance or rejection decision is shown in Figure 3.

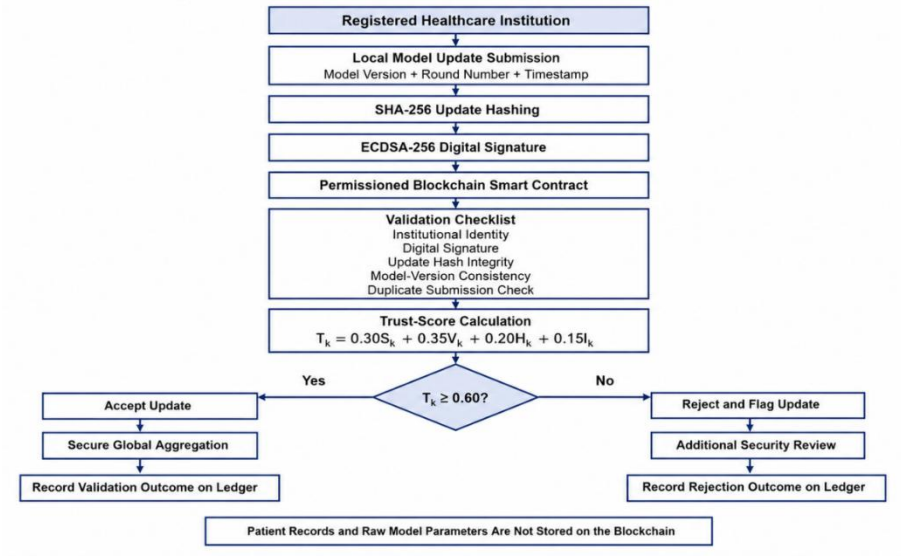

Figure 3. Blockchain-based model-update validation workflow**.**

4. EVALUATION PROTOCOL

4.1 Predictive and Comparative Evaluation

The proposed framework will undergo a comparison with four different models: LSTM in every institution, centralized LSTM, conventional Federated Averaging, and differentially private Federated Averaging. Installable LSTM will be modelled from institutions, leading to the performance achieved when healthcare institutions train independently. Centralized LSTM will be referenced for comparison only since it will need patient level records to be pooled. In conventional Federated Averaging, the benefits of collaborative learning will be defined without the consideration of further privacy safeguards; in differentially private Federated Averaging, the impact of the privacy-preserving noise will be included.

The input variables, the observation windows, the prediction horizon, and the partitioning of data and testing conditions will all be the same across all models providing a fair comparison. The full proposed framework will also include secure aggregation, blockchain verification and trust based on update filtering.

The accuracy, sensitivity, specificity, precision and F1-score will be used as measures of classification performance. PR-AUC will be given special attention in the event of imbalanced pandemic outcomes and ROC-AUC will be used to evaluate discrimination at classification thresholds. The Brier score and the expected calibration error will be used for assessing the reliability of probability predictions. Rounding errors will be used to check the 7-day forecasting for outbreaks, and these will be measured by mean absolute error and root mean squared error.

Individuals of ablation analysis will have one of the following removed: temporal LSTM modelling, differential privacy, secure aggregation, blockchain verification, and trust-based filtering. The methodological contribution of each framework component will be determined by the changes in predictive, privacy and security metrics.

4.2 Privacy, Security and Operational Evaluation

Differential-privacy budgets of ε = 0.5, 1, 2, 4, and 8 will be used to evaluate privacy protection. Membership-inference attacks will inform whether or not an intruder can uncover records that were used in the local training, and model-inversion attacks will evaluate if sensitive attributes can be reconstructed from model updates. The success of attacks will be contrasted in the different sets of the privacy budget to quantify the cost of privacy for the utility of the prediction.

Label flipping, random-gradient, model replacement and backdoor attacks will be part of the security evaluation. The attacks will be used at the following malicious participant rates: 10%, 20%, 30%. The evaluation will assess rates of poisoning detection, invalid-update rejection, and false-positive detection, success of the attack, and the degradation of the global model. As shown in Figure 4, the acceptance criteria are that the test should be positive at least 90% of the time, the test should reject updates which are invalid at least 95% of the time, and the test should succeed in membership-inference at most of the time, but not more that 55% of the time.

Blockchain transaction latency, transaction throughput, consensus time and validation success will be used as metrics for operational performance. The cost of communication will be based on the amount of megabytes sent out by each institution in each federated round. The scalability will be evaluated on different numbers of participating institutions (5 to 50) and will consider convergence time, total training time, blockchain overhead, and communication volume.

In reliability testing, communication failure will be added with peers, ordering-node, and institutional failures. The system has to function if 20% of the nodes involved in participation cannot be available. The main thresholds for the framework to meet the proposed validation requirements are summarized in figure 4, with each being either predictive, privacy, security, explainability or reliability.

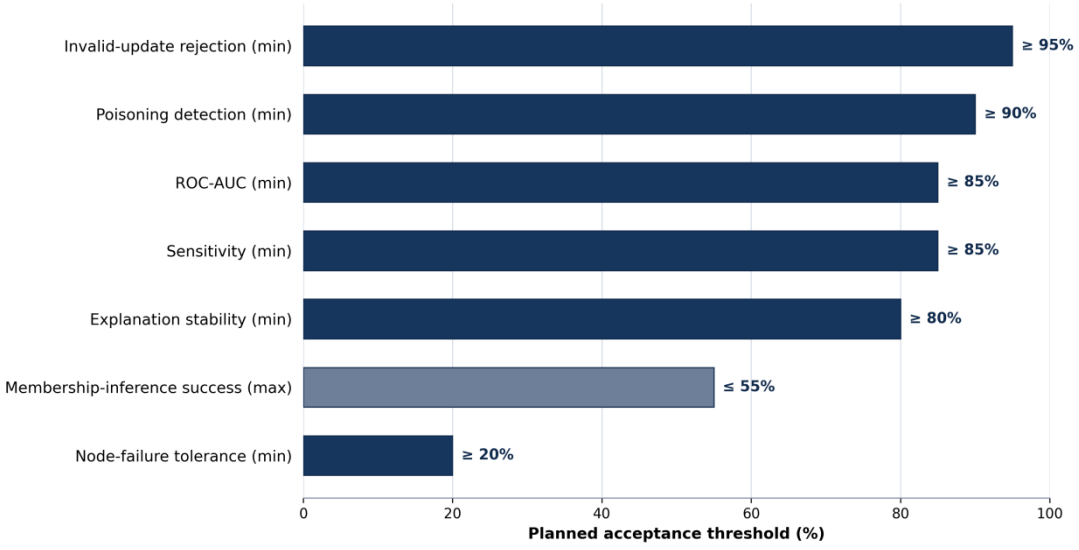

Figure 4. Performance and security acceptance criteria.

4.3 Fairness, Explainability and Statistical Analysis

It will be assessed based on pre-defined fairness, explainability and statistical acceptance criteria. F fairness analysis will assess the predictions across age of adolescent, sex, socioeconomic group, geographical location, adolescent vaccination and adolescent comorbidities. Each subgroup participants will have calculations of sensitivity, specificity, false-negative rates, and calibration. If there is a difference of more than 10 percentage points between groups, the data needs to be investigated, the model adjusted, or the data in the groups are to be calibrated separately.

Explanation stability and feature ranking consistency will be used to assess SHAP explanations. Whether similar records will get similar explanations in successive trials will be determined by explanation stability. Consistency in feature ranking will be used to assess if the most important features are consistent across different institutions, and across different versions of the models. Public-health workers will also evaluate the usefulness, clarity, and relevance of explanations generated.

The number of runs is set at 10, with different random seeds for each model configuration tested. Continuous measures will be displayed as a mean, standard deviation, and a 95% confidence interval. Proposed framework will be compared to comparison models using Wilcoxon sign rank test. A possible p value less than 0.05 will be considered statistically significant. When multiple models and subgroup outcomes are explored a multiple comparison correction will be used.

Table 3 shows the metrics and pre-defined criteria that are employed to assess predictive performance, privacy protection, security, operational efficiency, fairness, explainability and statistical reliability.

Table 3. Evaluation metrics and acceptance criteria

Dimension	Evaluation metric	Acceptance criterion
Predictive discrimination	ROC-AUC and PR-AUC	ROC-AUC ≥ 0.85
Classification	Sensitivity, specificity, precision and F1-score	Sensitivity ≥ 85%
Calibration	Brier score and expected calibration error	Calibration error ≤ 0.10
Forecasting	MAE and RMSE	Lower than persistence forecasting
Federated utility	Difference from centralized LSTM	Reduction ≤ 5 percentage points
Privacy protection	Membership-inference success rate	≤ 55%
Poisoning resistance	Malicious-update detection rate	≥ 90%
Update integrity	Invalid-update rejection rate	≥ 95%
Blockchain efficiency	Transaction latency and throughput	Latency ≤ 2 seconds
Communication efficiency	Data transmitted per round	≤ 5 MB per institution
Scalability	Convergence and training-time variation	Controlled increase from 5 to 50 institutions
Reliability	Operation under node failure	Operational under 20% node failure
Fairness	Maximum subgroup performance difference	≤ 10 percentage points
Explainability	SHAP explanation stability	≥ 0.80
Statistical reliability	Independent evaluation runs	Minimum 10 runs

5. DISCUSSION

The proposed architecture combines the key functionalities of federated temporal learning, differential privacy, secure aggregation, permissioned blockchain, trust-based update verification and explainable artificial intelligence in an integrated public-health surveillance solution. It is mainly methodological value for hospitals, laboratories and public-health agencies to cooperate without moving patient-level information down to a central repository. This is especially relevant when conducting surveillance of adolescents as the information may include clinical signs, vaccination status, exposure at school, household risk factors, socio-economic factors and geographical variables.

Federated learning allows for the aggregated health patterns to be distributed to the participating institutions without giving up data control. But alone local data retention does not ensure privacy. Model parameters could provide some metadata or attributes that could identify the membership. Differential privacy and secure aggregation can thus complement each other to create more comprehensive protections. Secure aggregation is a technique to ensure that the coordinating server does not get the individual institutional updates, while differential privacy restricts how much effect a single record has on the data. Lower privacy budgets would give better protection, but would limit the utility of the predictions; therefore, the privacy budget chosen should be traded off systematically with privacy–utility evaluation.

Permissioned blockchain offers institutional authentication, model-update integrity, tracking and tamper evident audit histories. However, blockchain should not be seen as a "privacy or cyber security fortress." It can prove if an update was sent from an authorized institution and it did not change, but it is unable to prove the validity of the update from a clinical perspective. Therefore, a trust based assessment is needed to detect label-flipping attacks, random-gradient attacks, model-replacement attacks, and backdoor attacks. It is important to also differentiate trustworthy updates from the institutions representing atypical populations from malicious contributions, so as to prevent the exclusion of underrepresented groups.

The minimum requirement are defined in the acceptance criteria given in Figure 4 and Table 3 to be assessed in the following step. Sensitivity and PR-AUC are especially noteworthy as results of pandemic surveillance could be sub-optimal and failing to identify an outbreak in a timely manner could result in false negative predictions. In addition to this, the reliability of risk probabilities should be taken into account as well as the ability to discriminate, for which calibration should be considered. The seven-day forecasting assessment will be based on the MAE and RMSE for the proposed temporal model, and compared to the simple persistence model.

The latency, communication amount, scalability and node-failure tolerance of a blockchain determine their operational feasibility. Enabling more and more institutions to participate (from 5 to 50) can improve the distributed learning but can also increase the amount of cryptographic processing, model-transfer, and coordination time. During institutional or network unavailability, the system will be evaluated under 20% node failure to see if the system can continue to operate.

Fairness analysis is necessary as performance of models can differ between different groups of adolescent girls (and boys) by sex, socioeconomic, geographical, vaccination, and comorbidities status. It is useful to take a look at subgroup sensitivity/caliberation rates and false negative rates, but overall rate is not enough. In the field of explanations, SHAP can show variables affecting prediction, but cannot be read as causal evidence, and rather refers to the statistical behavior of the model. Thus, the justification for the explanations will need to be judged, by public health workers, for its comprehensibility, relevance and actionability.

The framework is a method proposal that does not become a method until it is transferable to experiments and is evaluated. Ethics clearance will be necessary, adolescent data security measures, institutional governance arrangements, external validation, cyber security testing, fairness assessment and professional usability testing of its practical use will be required. This proposed architecture offers a structure for the future validation but is not proof in and of itself of effectiveness in clinical practice or readiness for a pandemic in the real world.

6. CONCLUSION

In this paper, we present a federated learning system based on blockchain for public-health surveillance and intelligence in a privacy-preserving setting, and specifically for health monitoring of adolescents. The two-layer LSTM, federated model aggregation, differential privacy, secure aggregation, permissioned blockchain verification, malicious update detection, and explanation using SHAP are integrated. It allows healthcare institutions to create models for surveillance together without centralizing patient-level data.

The permissioned blockchain offers institutional authentication, update-integrity verification, model provenance and an auditable validation history. Using differential privacy and secure aggregation decrease the risks that other participants might cause information disclosure, and trust-based filtering minimize the impact of unreliable or malicious participants. The framework facilitates 7 day infection-risk estimation, regional outbreak forecasting, cluster alerts and planning of healthcare resources.

A definition of a structured evaluation protocol has been formulated to evaluate predictive performance, calibration, privacy, security, communication efficiency, scalability, fairness, explainability and statistical reliability. These criteria define the

requirements for possible subsequent implementation, and not clinical outcomes. Future research needs to repeat case studies at multiple institutions, test for external validation, evaluate privacy and poisoning attacks, perform a subgroup-fairness study, measure professional usability, and then evaluate the study for prospective public-health study to ensure operational deployment.

REFERENCES

[1] Androulaki, Elli, et al. "Hyperledger fabric: a distributed operating system for permissioned blockchains." Proceedings of the thirteenth EuroSys conference. 2018.
[2] Bonawitz, Keith, et al. "Practical secure aggregation for privacy-preserving machine learning." proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. 2017.
[3] Choudhury, Atlanta, et al. "Leveraging federated learning and edge computing for pandemic-resilient healthcare." Scientific reports 15.1 (2025): 20497.
[4] Dayan, Ittai, et al. "Federated learning for predicting clinical outcomes in patients with COVID-19." Nature medicine 27.10 (2021): 1735-1743.
[5] Dwork, Cynthia, and Aaron Roth. "The algorithmic foundations of differential privacy." Foundations and trends® in theoretical computer science 9.3-4 (2014): 211-487.
[6] Kairouz, Peter, and H. Brendan McMahan. "Advances and open problems in federated learning." Foundations and trends in machine learning 14.1-2 (2021): 1-210.
[7] Kaissis, Georgios A., et al. "Secure, privacy-preserving and federated machine learning in medical imaging." Nature Machine Intelligence 2.6 (2020): 305-311.
[8] Kuo, Tsung-Ting, Rodney A. Gabriel, and Lucila Ohno-Machado. "Fair compute loads enabled by blockchain: sharing models by alternating client and server roles." Journal of the American Medical Informatics Association 26.5 (2019): 392-403.
[9] Lam, Maximilian, et al. "Gradient disaggregation: Breaking privacy in federated learning by reconstructing the user participant matrix." International Conference on Machine Learning. PMLR, 2021.
[10] Lundberg, Scott M., and Su-In Lee. "A unified approach to interpreting model predictions." Advances in neural information processing systems 30 (2017).
[11] McMahan, Brendan, et al. "Communication-efficient learning of deep networks from decentralized data." Artificial intelligence and statistics. Pmlr, 2017.
[12] Nguyen, Thien Duc, et al. "{FLAME}: Taming backdoors in federated learning." 31st USENIX security symposium (USENIX Security 22). 2022.
[13] Rieke, Nicola, et al. "The future of digital health with federated learning." NPJ digital medicine 3.1 (2020): 119.
[14] Sadilek, Adam, et al. "Privacy-first health research with federated learning." NPJ digital medicine 4.1 (2021): 132.
[15] Warnat-Herresthal, Stefanie, et al. "Swarm learning for decentralized and confidential clinical machine learning." Nature 594.7862 (2021): 265-270.