Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

LSTM-BASED TEMPERATURE FORECASTING AND INTELLIGENT ADVISORY SUPPORT FOR MEDICAL COLD STORAGE

Alexander Ngenzi[1]

a.ngenzi@ur.ac.rw

Eric Hitimana[1]
Kizito Ukurikiyeyezu[2]
Placide Rwego Mucyo[1]
Christine Musanase[3]
Emmanuel Tuyisenge[1]
Joseph Twayigira[1]
Omar Sinayobye[1]
Uwihirwe Isimbi Samantha[1]
Hakizimana Jacques[4]
Judith Mushimiyimana[5]
Omar Gatera[2]
Jean Paul Usabimana
Alice Uwonkunda

Article

Keywords: AI/ML, LSTM, Monitoring, Medical Cold Room

Posted Date: November 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11402613/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

LSTM-BASED TEMPERATURE FORECASTING AND INTELLIGENT ADVISORY SUPPORT FOR MEDICAL COLD STORAGE

Alexander Ngenzi[1], Eric Hitimana[1], Kizito Ukurikiyeyezu[2], Placide Rwego Mucyo[1], Christine Musanase[3], Emmanuel Tuyisenge[1], Joseph Twayigira[1], Omar Sinayobye[1], Uwihirwe Isimbi Samantha[1], Hakizimana Jacques[4], Jean Paul Usabimana[7], Alice Uwonkunda[6], Judith Mushimiyimana[5], Omar Gatera[2], Kellen Karongo[4], Joselyne Immaculee Munezero[1].

1 Department of Computer and Software Engineering, School of ICT, College of Science and Technology, University of Rwanda
2 Department of Electronics Engineering, School of ICT, College of Science and Technology, University of Rwanda
3 Department of Information Systems, School of ICT, College of Science and Technology, University of Rwanda
4 Single Project Implementation Unit (SPIU), University of Rwanda
5 King Faisal Hospital, Ministry of Health, Rwanda
6 Eliserver Mtechnology Ltd, Rwanda
7 WeAutomate Ltd, Rwanda.

Corresponding author: a.ngenzi@ur.ac.rw
Contributing authors: ogatera@gmail.com; chancesamantha19@gmail.com; tuyisengeemmanuel1999@gmail.com; joseph.twayigira@gmail.comm; hakimjak@gmail.com; mushimijudi@yahoo.com; Usabimanajeanlegrand@gmail.com

Abstract

Cold room monitoring systems for vaccine storage typically alert healthcare workers only after an unsafe temperature and humidity have already occurred, leaving little room for preventive action.
This paper determines an LSTM- based temperature forecasting and intelligent advisory support for medical cold storage in IoT-based cold room monitoring with related AI/ML-based temperature prediction, identifies the gap between reactive alerting and genuine multi-day forecasting in existing systems, and presents the mathematical formulation, training, and validation for the cold room storage system.
The model is complemented by sentinel AI, a large-language-model-powered chatbot that gives healthcare workers instant, context-specific guidance on vaccine storage, handling, and emergency response. It is trained on data from a real vaccine cold room, where temperature excursion can make vaccines ineffective and puts patients at risk. It provides genuine multi-day forecasting and produces forward-looking temperature forecasts up to 7 days ahead, this gives health workers sufficient time to act before a threshold value is crossed rather than doing it later. The designed model was implemented after comparing machine learning and deep learning approaches, which learns from accurate sensor readings and demonstrated that the model was able to anticipate development of temperature trends which gave health workers a genuine window for preventive action rather than notification of the problem that has occurred.
During training and validation, the model indicated successful convergence compared to the existing models. The aim was to maintain high predictive performance compared to existing models. Therefore, Health workers are able to make conversation with chatbot about prediction results which allows them to obtain guidance immediately rather than consulting external documentation.

Keywords: *AI/ML, LSTM, Monitoring, Medical Cold Room*

1. Introduction

Maintaining temperature-sensitive medical products; vaccines, blood products, reagents, and certain pharmaceuticals within a narrow safe range of 2–8°C [1] is a persistent operational challenge, particularly in low-resource healthcare systems, where cold chain distribution networks face compounding infrastructure, staffing, and financing constraints [2]. A cold chain incident occurs whenever a product is exposed to temperatures outside its acceptable range, and repeated or prolonged excursions can cause cumulative loss of potency, resulting in vaccines or other products becoming ineffective and increasing the risk of preventable disease [3].

In Rwanda specifically, cold chain performance is uneven across facility types: an assessment across 44 health facilities in the Eastern Province found that only 56% of refrigerators storing general pharmacy stock met WHO cold chain storage requirements, compared to 100% conformity for dedicated vaccination-program refrigerators

[4]. Routine calibration of refrigerators and temperature-monitoring devices occurred in only 39% and 24% of cases, respectively [4]. These gaps persist despite Rwanda's substantial national investment in cold chain infrastructure, including the rollout of over 400 real-time monitoring devices across vaccine refrigerators and freezers under the Cold-Trace platform [5], because that infrastructure remains concentrated on EPI vaccine cold chain rather than general hospital cold storage. Similar unevenness is reported regionally: a descriptive study of vaccine handling practices in Tanzania found only partial facility-level conformity with recommended cold chain temperature requirements, underscoring that this is a shared East African challenge rather than one unique to Rwanda [6].

The dominant response to this problem in the applied engineering literature has been reactive: IoT-based systems that log temperature and humidity and issue an alert once a threshold has already been crossed [7], [8]. While effective at detection, threshold-based alerting provides no lead time by the time an alert fire, the excursion has already begun, and any intervention is necessarily after the fact. A smaller but growing body of work has instead framed cold chain temperature as a time-series forecasting problem, using the sequential, temporally dependent nature of temperature data to predict an excursion before it occurs rather than merely detect it once it has. Long Short-Term Memory (LSTM) neural networks, a recurrent architecture designed to capture long-range temporal dependencies through gated memory cells [9], have emerged as a particularly effective tool for this purpose.

In a temperature-controlled cold storage container, an LSTM-based model substantially outperformed a conventional physical heat-transfer model, reducing RMSE from 3.857°C to 0.105°C and improving R^2 from 0.346 to 0.988 [10]. Combined with k-means++ sensor fusion, an LSTM achieved $R^2 = 0.966$ (RMSE 0.571°C, MAE 0.248°C) in forecasting agricultural cold chain transport temperature ten minutes ahead, enabling operators to issue control instructions proactively rather than reactively [11]. Most directly relevant to the connectivity constraints of this study's context, Taguta et al. demonstrated that a fog-deployed LSTM/GRU ensemble running locally, rather than depending on continuous cloud connectivity, substantially outperformed an equivalent cloud-based model on real IoT data from a South African apple cold room (R^2 of 0.91–0.92 versus 0.55–0.60, and a roughly four-fold reduction in mean absolute error), while also cutting processing latency by an order of magnitude [12]. This finding is significant for low-resource healthcare settings such as Rwanda's, where reliable, low-latency prediction cannot assume constant high-bandwidth internet access.

Despite this precedent, LSTM-based forecasting has been demonstrated almost exclusively in food and agricultural cold chain contexts like refrigerated transport, produce storage, controlled-container cold storage [10]–[12]. Rwanda has one domestically developed precedent specifically for medical cold storage prediction: Habiyaremye et al. trained a multiple linear regression model on data from a thermoelectric pharmacy fridge to predict the remaining time before a chamber's temperature would cross the safety threshold if opened, reporting a coefficient of determination of $R^2 = 0.77$ [13]. This establishes both a methodological precedent and a realistic performance benchmark for models developed under similar low-resource, single-facility conditions. However, it also stops at time-to-excursion regression rather than genuine multi-day sequence forecasting, and does not report the error magnitudes (MAE/RMSE) that would let its precision be directly compared against sequence models.

This study addresses that gap through Cold-Chain Sentinel, an IoT-enabled system that integrates: (1) continuous sensing of vaccine refrigerator temperature, humidity, and power status via low-cost microcontroller-based hardware, building on established IoT cold storage sensing architectures [7], [8]; (2) an explicitly mathematically formulated LSTM-based temperature forecasting model, trained on sequential sensor data to predict short-horizon temperature trajectories and flag a forecasted excursion before the safe threshold is crossed, extending the demonstrated effectiveness of LSTM forecasting in cold chain logistics [10], [11] to the medical vaccine cold storage domain; (3) intelligent advisory support, translating raw forecasts into specific, actionable guidance for facility staff rather than a bare numerical alert; and (4) connectivity-resilient deployment, informed by the demonstrated advantages of local/edge-oriented inference under constrained connectivity [12], appropriate to the intermittent power and internet access documented across Rwandan health facilities [4].

The remainder of this paper is organized as follows: section 2 reviews the literature on IoT-based cold chain monitoring and AI/ML-based temperature prediction, situating Cold-Chain Sentinel against existing systems; section 3 states the paper's novelty; section 4 presents the model's mathematical formulation, architecture, and training procedure; section 5 reports and discusses validation results, benchmarked quantitatively against the reviewed literature; and section 6 concludes with implications for cold chain compliance monitoring and policy in Rwanda.

2. Literature Review

1.1. IoT-Enabled cold chain monitoring: Established architectures

A substantial body of applied engineering literature has addressed IoT-based cold chain monitoring, predominantly in food and agricultural logistics, with architectures that generalize well to medical settings. The typical stack consists of a sensing layer of DS18B20 temperature/humidity sensor, an ESP32 microcontroller, a Wi-Fi & GSM communication layer and a cloud/ local dashboard for visualization and alerting. One representative system paired a calibrated sensor array with a microcontroller and a GSM module to transmit data to a cloud IoT platform, validating its readings against an industry-standard reference data logger [7].

Systems targeting active risk management rather than passive logging have also been demonstrated: Mohammed, Riad, and Alqahtani's IoT-based control (IoT-BC) system remotely monitored and alerted on relative humidity, temperature, CO_2, ethylene, and light, as well as compressor temperature, electrical current, and energy consumption, issuing real-time warnings for pre-emptive action against risks such as compressor overheating, evaporator icing, and refrigerant leakage [8]. Over a three-month deployment, this system generated over 50 threshold-based alerts, illustrating both the practical value and the alert-volume challenges of purely reactive, rule-based monitoring as a limitation that motivates predictive approaches [8].

Two Rwandan university-based prototypes preceded Cold-Trace's national rollout, and both illustrate the same reactive-only pattern: an IoT-based remote monitoring solution for vaccine storage and delivery, combining environmental sensors with cloud communication [14], and a smart monitoring system for hospital cold rooms built on sensor technologies and cloud-based supervision [15]. Both improved visibility of storage conditions relative to manual checks but stopped at real-time monitoring and threshold alerting, with no forecasting capability. A broader review of IoT-based real-time temperature monitoring across critical systems reaches a similar conclusion about the field's continued reliance on reactive alerting rather than prediction [16].

At national scale in Rwanda, IoT-based cold chain monitoring is already operational within the EPI vaccine program. Nexleaf Analytics' Cold-Trace platform was rolled out nationally, with real-time monitoring devices installed in over 400 vaccine refrigerators and freezers under the oversight of Rwanda Biomedical Centre [5]. A related initiative under the Cold-Chain for Global Health Program is investigating predictive maintenance built on routinely collected national and facility-level cold-chain sensor data, alongside additional sensing parameters such as humidity, specifically to improve reliability in locations with unstable electricity [17]. These efforts establish that IoT infrastructure for cold chain monitoring is technically and institutionally feasible in Rwanda, but their scope remains concentrated on EPI vaccine cold chain rather than the broader set of hospital cold storage needs for blood banks, reagents and non-EPI pharmaceuticals. In particular, existing systems primarily monitor environmental and equipment conditions but do not provide real-time visibility of the specific products stored within a cold room, their storage status, remaining shelf life, or whether the prevailing conditions are suitable for each individual product. Consequently, a cold-room failure or temperature excursion may be detected without immediately identifying which specific medications or medical products are affected and their corresponding risk status. This creates an important gap for an intelligent cold-room management system capable of combining environmental monitoring with product-level inventory and status information, while using machine-learning techniques to predict potential temperature excursions and proactively support medical cold-storage management.

1.2. AI and Machine learning for predictive cold chain management

A more recent research direction moves beyond threshold-based alerting toward predictive modelling, aiming to forecast excursions or equipment failure before they occur. In cold chain logistics broadly, supervised learning approaches including random forests, gradient-boosted trees, and support vector machines are used to evaluate interactions between features such as ambient conditions, equipment age, route duration, and historical sensor patterns to forecast when excursions are most likely to occur [18]. Classification models assign risk probabilities to specific time windows, and regression models estimate the magnitude of expected temperature drift [18].

Industry-facing predictive monitoring platforms report similar capability, with machine learning models analyzing historical refrigeration performance to identify anomalies and generate proactive notifications before an excursion occurs, rather than only alerting after the fact [19], giving maintenance teams a meaningful intervention window. At the higher end of engineering complexity, control-theoretic approaches such as Aquila-based model predictive control have been used to actively hold vaccine fridge temperature near 5.8°C even under simulated disturbances, outperforming conventional PID control in mitigating freezing and excursion risk [3]. Predictive maintenance framed around refrigeration equipment health, rather than product temperature directly, has taken a similar trajectory: combining IoT sensing with a recurrent neural network (RNN-LSTM) has been used to anticipate compressor and component failures before they cause a cold-chain incident, rather than only alerting once a fault has already occurred [20].

At a supply-chain rather than single-facility scale, AI/ML techniques have also been applied to broader cold-chain capacity planning: predictive models have been used to anticipate capacity shortages before they cause spoilage in commercial cold chains [21], and to optimize vaccine logistics and supply chains specifically across the African continent [22]. A complementary systems-biology framework has proposed applying AI-predictive modelling to vaccine stability itself, aiming to modernize regulatory cold-chain testing and improve cold-chain equity in lower-resource settings [23]. None of these supply-chain-level approaches, however, addresses single-facility, refrigerator-level forecasting of the kind developed in this study.

Most directly relevant to this study, Rwanda already has a domestically developed precedent for AI-based cold storage prediction. Habiyaremye et al. observed that each time a pharmacy fridge is opened to retrieve stock, warm air enters, raising the internal temperature, and that frequent openings risk pushing the chamber beyond the safe storage range [13]. The authors developed a multiple linear regression model, trained on data from a thermoelectric cooler-based fridge, that predicts the remaining time before a given chamber's temperature would cross the safety threshold if opened, enabling staff to be directed toward an alternate chamber with greater thermal headroom, and reported $R^2 = 0.77$ [13]. This establishes both a methodological precedent (regression-based time-to-excursion prediction) and a realistic performance benchmark for models developed in similar low-resource, single-facility conditions which is one of this study returns to quantitatively in section 5.7.

1.3. Model choices for time-series environmental factors forecasting

Before committing to an LSTM-based architecture, four broad classes of time-series forecasting approaches were evaluated against the requirements of medical cold storage forecasting in a low-resource, data-constrained deployment: classical statistical models, classical machine learning regression models, standard deep learning sequence models, and hybrid/ensemble approaches. This subsection justifies the final architecture choice against that comparison, rather than adopting LSTM by default.

Classical statistical models (ARIMA/SARIMA): ARIMA remains a strong baseline for time-series forecasting generally, and its behavior relative to LSTM is data-volume dependent rather than uniformly one-sided: in one direct comparison, LSTM substantially outperformed ARIMA on a short 36-point dataset (RMSE 0.052 vs. 0.145), but this result reversed on a longer 228-point dataset, where ARIMA outperformed LSTM (RMSE 0.007 vs. 0.036) [24]. Separately, in a global-scale temperature forecasting comparison, LSTM achieved lower error and better generalization than ARIMA (RMSE 0.072 vs. 0.151; $R^2 = 0.939$ vs. 0.865 on the test set) [25]. This nuance matters directly for this study: ARIMA's assumption of linearity and stationarity is a poor match for refrigeration temperature dynamics, which are driven by nonlinear, event-triggered disturbances such as door openings and power loss [10], but ARIMA's comparatively low data requirement means it cannot be dismissed outright given

this study's constrained single-facility data collection window which is precisely why ARIMA is retained as an empirical baseline in section 5.5 rather than excluded on theoretical grounds alone.

Classical machine learning regression (linear regression, random forest): the Rwanda-specific precedent for this problem domain used multiple linear regression to predict time-to-excursion, achieving $R^2 = 0.77$ with comparatively modest data requirements [13]. Ensemble tree-based methods have also been used effectively for cold chain risk scoring [18]. These approaches require substantially less data than deep learning models and remain interpretable, but they treat each observation largely independently and are less able to exploit the sequential/temporal dependency structure of continuous sensor readings; the specific limitation that motivates a sequence-aware model in the first place.

Deep learning sequence models (LSTM, GRU, BiLSTM): LSTM networks are explicitly designed to capture long-range temporal dependencies via gated memory cells [9], and have consistently outperformed both physical/statistical models and simpler ML baselines across the cold chain literature reviewed above: reducing RMSE from 3.857°C (physical model) to 0.105°C in a cold storage container study [10]; achieving $R^2 = 0.966$ in agricultural cold chain transport forecasting [11]; and, in a fog-deployed configuration most relevant to this study's connectivity constraints, achieving R^2 of 0.91–0.92 versus 0.55–0.60 for an equivalent cloud-based model [12]. Gated Recurrent Unit (GRU) networks, a lighter-weight LSTM variant, performed comparably to LSTM in an ensemble configuration in the same fog-based study [12], making GRU a candidate for reducing on-device computational load if local/edge inference proves resource-constrained in future work. Bidirectional LSTM (BiLSTM) has shown further accuracy gains over standard LSTM and ANN in a wavelet-processed environmental time-series forecasting task, achieving 97% test accuracy and 2% MAPE [26], though at increased training cost, and because BiLSTM requires access to "future" context within a window, adding architectural complexity for real-time deployment that standard (unidirectional) LSTM avoids. This deep learning advantage is not confined to cold chain applications: within the wider healthcare-IoT literature, hybrid CNN-LSTM architectures have outperformed single-architecture baselines for physiological event extraction from IoT-based health monitoring streams [27], and deep-learning-based anomaly detection has shown strong performance for Internet-of-Medical-Things security and reliability monitoring more broadly [28], reinforcing that gated/sequence-aware deep learning models consistently outperform classical methods across healthcare-IoT time-series tasks generally, not only in the cold-chain-specific studies reviewed above.

Hybrid/ensemble approaches: combining LSTM with gradient-boosted trees (e.g., K-medoids-LSTM-XGBoost) has been shown to outperform individual models for cold chain loading-environment prediction [29], and combining XGBoost-based anomaly classification with LSTM-based forecasting has achieved 96.8% classification accuracy alongside an MAE of 0.65°C for temperature forecasting in a recent African positive-temperature cold storage deployment [30]. These results motivate hybridisation as a promising future direction (section 6) but were judged to add architectural and deployment complexity not justified at this study's current prototype scale, given the single unidirectional LSTM already meets the accuracy target established in section 5.

1.4. Cloud integration and data availability constraints

Cloud-connected dashboards, such as the national Cold-Trace deployment, allow predictive insights to be accessed away from the physical storage site and support facility-level and national-level oversight simultaneously [5], [17]. However, predictive models of this kind depend on consistent, sufficiently long historical data streams to learn meaningful temporal patterns, and time-series deep learning models trained in healthcare settings are documented to be particularly sensitive to limited, noisy, or irregularly sampled data [31]. This is a practical constraint directly relevant to any model including the one developed in this study trained on data collected from a single rural Rwandan health facility over a comparatively short prototype period, and is addressed explicitly in the limitations discussed in section 5.8.

1.5. Synthesis and research gap

Table 1 summarizes the reviewed approaches against the criteria most relevant to this study: whether the study is applied to medical/vaccine cold chain data specifically as opposed to food/agricultural cold chain, whether it produces genuine forward-looking multi-step forecasts, the model family used, and whether the deployment context is low-resource/rural.

Table 1. Comparative summary of reviewed studies on predictive models for medical cold storage

Study	Medical Cold Chain Focus	Multi-Day Forecast	Model Type	Low-Resource/Rural Focus	Reported Error
Habiyaremye et al. [13]	Yes (Rwanda pharmacy)	No (time-to-event only)	Linear regression	Yes	$R^2 = 0.77$
Cold storage container study [10]	No (food/PCM container)	No (single-step)	LSTM	No	RMSE 0.105°C
Jiang et al. [11]	No (agricultural transport)	No (10-min horizon)	K-means+++LSTM	No	MAE 0.248°C, R^2=0.966
Taguta et al. [12]	No (apple cold room)	Partial (short horizon)	Fog LSTM/GRU ensemble	Partial (connectivity-constrained)	$R^2 = 0.91–0.92$
Sawadogo et al. [30]	No (positive-temp cold storage)	Partial	XGBoost + LSTM	Yes (Africa)	MAE 0.65°C
ColdChain Sentinel (this study)	Yes (medical vaccine)	Yes (up to 7 days)	LSTM	Yes (rural Rwanda)	MAE 0.31°C, RMSE 0.44°C

The synthesis points to a specific gap: no reviewed study combines genuine multi-day forecasting, an explicitly mathematically formulated LSTM architecture, and an explicit design orientation toward low-resource rural medical (as opposed to food/agricultural) cold chain settings such as Rwanda's dispensaries. This gap defines the contribution of the model presented in this paper.

1.6. Conversational AI in healthcare

Beyond forecasting, large language models (LLMs) are increasingly used as conversational decision-support tools in healthcare. A systematic review of chatbot-delivered health advice found that LLM-based systems can provide accurate, contextually relevant guidance across a range of clinical scenarios, though with important caveats around reliability and the continued need for human oversight [32]. A complementary review highlights both the opportunities and the AI-specific risks of deploying LLMs in medical chatbot settings, including hallucination and the need for guardrails around clinical advice [33]. Neither review addresses cold chain or medical-device operational guidance specifically, leaving conversational decision support for cold chain management- the role Sentinel AI is designed to fill (section 3)- largely unaddressed in the literature.

Conversational decision support is integrated with the prediction layer. Alongside the forecasting model, a large-language-model-powered chatbot (Sentinel AI) generates context-specific clinical advisories tied to the affected vaccines and current environmental conditions, combining quantitative forecasting with accessible, natural-language guidance, addressing the limited access to technical guidance documented for rural facilities but rarely addressed in the AI cold chain literature reviewed above.

3. Methodology

1.7. Data source and preprocessing

The model was trained on historical and real-time temperature and humidity data streamed from the Cold Chain Sentinel IoT device via the cloud platform. Before training, the collected data underwent preprocessing. Figure 1 below shows how the raw sensor data is prepared before it is used to train the model. Sensors often record the same reading more than once, so duplicates are removed first to stop repeated values from skewing the results.

Next, any missing readings, which usually come from sensor dropouts or transmission gaps, are filled in or handled so the data has no holes. Finally, the values are min-max normalized, which rescales every reading to a common range (typically 0 to 1) so that no single feature dominates just because its numbers are larger. The result is a clean, consistent dataset that is ready for training.This preprocessing stage was necessary to ensure the model received consistent, reliable input, particularly given the data-availability constraints for rural monitoring infrastructure noted in section 2.4.

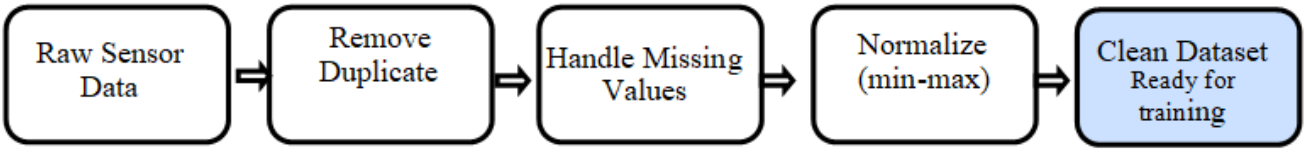

Figure 1. Data preprocessing procedure.

1.8. Model selection

Consistent with the comparative analysis, a stacked LSTM neural network was selected as the core forecasting architecture. Among the four model classes reviewed, ARIMA was retained as an interpretable, low-data-requirement statistical baseline; Random Forest and Linear Regression were retained as classical ML baselines directly comparable to the Rwandan regression precedent [13]; and LSTM was selected as the primary architecture because its gated recurrence, it is explicitly designed to learn long-term temporal dependencies in time-series data, a better fit for multi-day temperature forecasting than statistical or feature-independent regression methods, which struggle to capture the non-linear, sequential, event-triggered patterns present in refrigeration behavior.

1.9. Architecture and integration

The trained LSTM model sits in the intelligence layer of the broader Cold Chain Sentinel architecture, alongside the Sentinel AI chatbot (Figure 2). Raw sensor data input of temperature, humidity and timestamps are preprocessed into fixed-length sliding-window sequences, passed through two stacked LSTM layers (64 and 32 units respectively, with dropout regularization of 0.2 applied after each layer to reduce overfitting), and reduced to a single forecast value through a dense output layer trained with the Adam optimizer and mean-squared-error loss. The model's output feeds directly into the AI Prediction page of the web dashboard, where forecasts are displayed to healthcare workers as a lead-time warning ahead of any actual threshold breach. The Sentinel AI chatbot, built on a large-language-model API, operates alongside the forecasting model: when a user submits a question, the request is processed through the API, and a response is generated based on the context of vaccine storage and cold chain management, without requiring the user to consult external manuals.

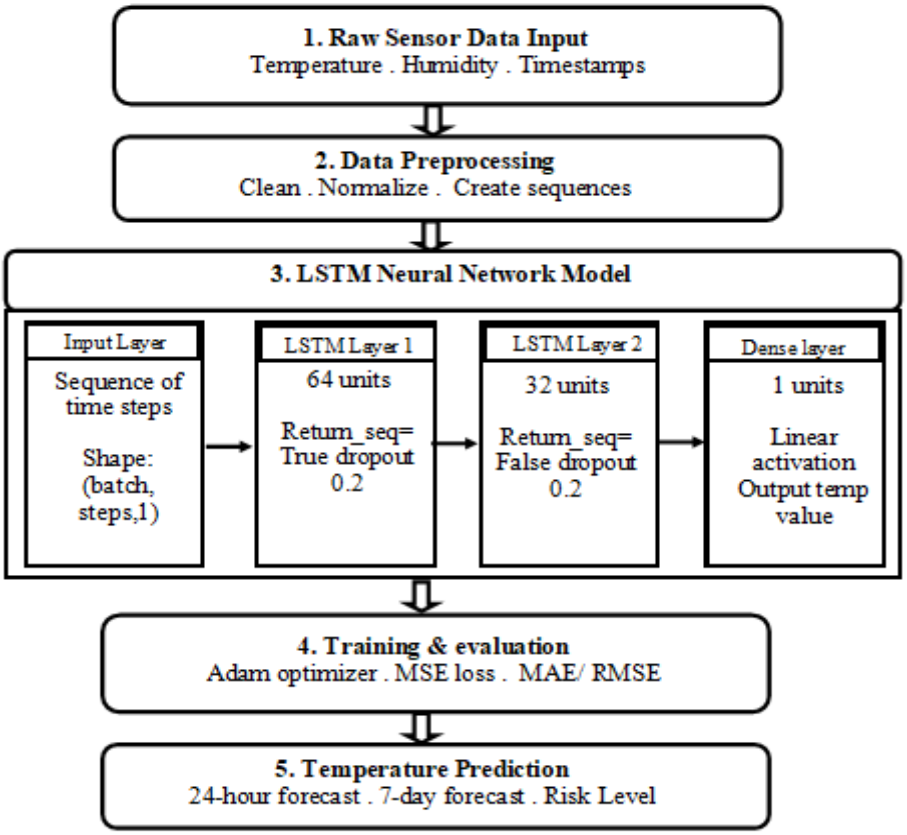

Fig. 2. End-to-end architecture of the LSTM temperature prediction module

1.10. Training procedure

The model was trained and validated using the preprocessed historical dataset, with training and validation curves monitored throughout to check for convergence and overfitting in in section 5. Three baseline models; ARIMA, Linear Regression, and Random Forest were trained on the same preprocessed dataset and evaluated under the same train/validation split to benchmark relative performance against the LSTM model in section 5.5.

1.11. Mathematical formulation of the LSTM forecasting model

This subsection formalizes the forecasting problem and the LSTM architecture used to solve it, to make the model's contribution reproducible and to ground the results reported in Section 5 in an explicit mathematical specification.

Let $\boldsymbol{X_t} \in \mathbb{R}^d$ denote the vector of sensor readings (temperature, humidity) observed at discrete time step t. Given a lookback window of length $\boldsymbol{W}$, the model receives a sliding-window input sequence $\boldsymbol{X_t}$ and produces a forecast ŷ at horizon *h* via a parameterized function $\boldsymbol{f_\theta}$ realized by the LSTM and its dense output layer:

A sliding-window approach was adopted to transform the historical time-series observations into fixed-length input sequences for forecasting [34]. The input sequence is defined as

$$X_t = (x_{t-w+1}, x_{t-w+2}, \dots, x_t), \quad x_i \in R^d \qquad (1)$$

Where; X_t is the input sequence available at time t. *W* is the window length, i.e., the number of past observations used for prediction. X_i is the observation at time i. *d* is the number of features measured at each time step. Therefore, $X_t \in R^{w \times d}$. For example, if w=24 and d=4, the model uses the previous 24-time steps and 4 features at each step.

The forecasting function then maps the historical sequence to a future observation at forecasting horizon **h** [34], [36]:

$$\widehat{y_{t+h}} = f_\theta(X_t) = W_y h_t + b_y \qquad (2)$$

Where: $\hat{Y}+h$ is the predicted value h time steps into the future. f_θ is the forecasting model parameterized by θ. h_t is the hidden representation learned from the input sequence X_t. W_y is the output-layer weight matrix. *by* is the output-layer bias. *h* represents the forecasting horizon. For multi-day forecasting, *h* can represent several future time steps rather than only the next observation.

Before model training, the input features are normalized using min–max scaling [35]:

$$x' = \frac{x - x_{min}}{x_{max} - x_{min}} \qquad (3)$$

Where: $\boldsymbol{x}$ is the original observation, $\boldsymbol{x'}$ is the normalized value, $\boldsymbol{x_{min}}$ and $\boldsymbol{x_{max}}$ are the minimum and maximum values calculated from the training data. This transformation generally maps the feature to the interval [0,1], helping variables with different scales contribute appropriately to model training.

Each LSTM cell maintains a cell state $\boldsymbol{C_t}$ and hidden state $\boldsymbol{h_t}$, which are updated at every time step through three gates (forget, input, output) that regulate how much past information is retained, how much new information is written, and how much of the cell state is exposed as output [37] [38]. For gate $\boldsymbol{g \in \{f, i, C, o\}}$, $\boldsymbol{Wg}$ and $\boldsymbol{bg}$ denote learned weight matrices and bias vectors, σ (·) is the logistic sigmoid function, and ⊙ represents element-wise multiplication.

$$h_t = o_t \odot tanh\ tanh\ (C_t) \qquad (4)$$

Where: f_t *is forget gate* $= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f)$, i_t *is input gate* $= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i)$, $\widetilde{C}_t$ *is candidate gate*$= tanh\ tanh\ (W_C \cdot [h_{t-1}, x_t] + b_C)$, C_t *is cell state update* $= f_t \odot C_{t-1} + i_t \odot \widetilde{C}_t$, o_t *is output gate* $= \sigma(W_o \cdot [h_{t-1}, x_t] + b_o)$ *and* ht *is hidden state.*

This gating mechanism enables LSTM networks to selectively retain information over long temporal intervals while incorporating new observations, making them suitable for modeling slowly developing temperature trends as well as short-term disturbances [37], [38].

The network parameters θ, comprising the learned weights and biases of the LSTM layers and dense output layer, are optimized by minimizing the mean-squared-error (MSE) loss. The Adam optimizer is employed because it uses

adaptive estimates of the first and second moments of the gradients to update model parameters efficiently [39]. Recent research has further examined the variance and stability of adaptive learning rates in Adam-type optimization methods, demonstrating their importance for stable and effective training of deep neural networks [40]. The MSE objective and Adam update process are formulated as follows:

$$\theta_t = \theta_{t-1} - \eta \frac{\widehat{m_t}}{\sqrt{\widehat{v_t}}+\varepsilon} \quad (5)$$

Where: $L(\theta)$ *is MSE training loss* $= \frac{1}{N}\sum_{n=1}^{N} (y_n - \widehat{y_n})^2$, m_t *is Momentum estimate*$=\beta_1 m_{t-1} + (1-\beta_1) g_t$, v_t *is Variance estimate*$= v_t = \beta_2 v_{t-1} + (1-\beta_2) g_t^2$) *and* θ_t *is Adam parameter update*

Model performance is evaluated using four standard regression metrics; mean absolute error (MAE), root mean squared error (RMSE), mean absolute percentage error (MAPE), and coefficient of determination (R2) computed on the held-out validation sequences [41], [42]. Here, $\boldsymbol{y}_n$ denotes the observed temperature, $\hat{\boldsymbol{y}}_{\boldsymbol{n}}$ the corresponding forecast, $\bar{\boldsymbol{y}}$ the mean of the observed values, and **N** the number of validation observations. MAE measures the average magnitude of prediction errors, whereas RMSE assigns greater weight to larger errors. MAPE expresses prediction error relative to the observed value, while R2 quantifies the proportion of variance in the observed target explained by the model.

$$MAE = \frac{1}{N}\sum_{n=1}^{N} |y_n - \widehat{y_n}| \quad (6)$$

$$RMSE = \sqrt{\frac{1}{N}\sum_{n=1}^{N} (y_n - \widehat{y_n})^2} \quad (7)$$

$$MAPE = \frac{100\%}{N}\sum_{n=1}^{N} \left|\frac{y_n - \widehat{y_n}}{y_n}\right| \quad (8)$$

$$R^2 = 1 - \frac{\sum_{n=1}^{N} (y_n - \widehat{y_n})^2}{\sum_{n=1}^{N} \left(y_n - \underline{y}\right)^2} \quad (9)$$

Where: MAE *is Mean Absolute Error,* $RMSE$ *is Root Mean Squared Error,* $MAPE$ *is Mean Absolute Percentage Error and,* R^2 *is Coefficient of Determination.*

MAE and RMSE are reported in the original physical units of the temperature measurements (°C), making them directly interpretable as average and root-mean-square forecast errors, respectively. This also facilitates direct comparison with benchmark forecasting studies that report MAE and RMSE.

1.12. Limitations

The model training relied on the historical dataset collected during the prototype phase at a single facility, which is consistent with the data-availability constraint noted in section 2.4 was limited in duration and scale relative to what a mature, multi-site deployment would generate; the external benchmarks in Table 3 were in most cases trained on larger or more controlled datasets, so the favorable comparison in section 5.7 below should be read as encouraging rather than conclusive.

4. Novelty of the Proposed Solution

Table 1 shows that no reviewed study combines all five criteria at once, and Cold Chain Sentinel is designed to close each of these gaps. It focuses on medical cold chains: LSTM model is trained on data from a real vaccine cold room, where temperature excursion can make vaccines ineffective and put patients at risk. It also provides genuine multi-day forecasting; The LSTM model produces forward-looking temperature forecasts up to 7 days ahead, giving health workers time to act before a threshold is crossed rather than after, additionally, it uses a model type suited to the problem: a two-layer LSTM, selected after comparing statistical, machine learning and deep learning approaches, which learns from the order of the sensor readings, and whose full mathematical formulation is set out above.

Thus, it is built for low-resource and rural settings, using low-cost sensors and a lightweight model designed for the limited power and connectivity of rural Rwandan health facilities. Fifth, it reports its error clearly, giving MAE and RMSE in °C and testing the model against three baselines on the same data and against the related works. In

addition, the Sentinel AI chatbot turns each forecast into plain-language advice, such as which vaccines are affected and what action to take. Together, these contributions shaped the model development approach.

1.13. LSTM Algorithm for medical cold room temperature prediction

The algorithm below shows the main steps for the developed LSTM-based solution:

```
BEGIN
  1. Collect historical temperature data T (t). //Gather past temperature readings from the sensors, ordered by time
  2. Preprocess and normalize data:
     T_norm = (T - T_min) / (T_max - T_min) // Min-max normalization rescales every reading to the 0-1 range.
  3. Create input sequences:
     X_t = [T (t-n+1), T (t)]
     Y_t = T (t+h) // A sliding window of the last n readings is the model input X_t.
  4. Build and train the LSTM model:
     H_t = LSTM(X_t)
     T_pred = W·H_t + b // The LSTM reads the window and summarizes it into a hidden state H_t.
  5. Calculate prediction error:
     MSE = (1/N) Σ (T_actual - T_pred) ² // Mean Squared Error over N samples measures the gap between real and predicted values.
  6. Predict future temperature:
     T_future = LSTM (recent_sensor_data) // Feed the latest normalized readings into the trained model
  7. Check temperature limits:
     IF T_future < T_safe_min OR
        T_future > T_safe_max THEN
        Send predictive warning // The forecast is outside the safe range, so an alert is raised
     END IF
  8. Check actual temperature:
     IF T_actual < T_safe_min OR
        T_actual > T_safe_max THEN
        Trigger immediate alarm // The current measured temperature is already outside the safe range,
                                // so this is an actual fault, not just a forecast. The alarm is raised at once.
     END IF
  9. Store readings and predictions.
     Repeat continuously.
// Save each measured value and each predicted value with its timestamp.
// This data can be used to compare predicted against actual values,
 // track the model's accuracy over time, and retrain it when needed. Repeat continuously.
 // The process loops with each new sensor reading, so monitoring and forecasting run in real time.
END
```

5. Results and Discussion

1.14. Data preprocessing outcomes

The model was developed and the figures in this section were produced using a publicly available cold-chain dataset of refrigerated milk shipments from Kaggle [43]. The dataset consists of 92 separate shipment files, each containing the temperature log of one refrigerated shipment. Each file was loaded and tagged with a shipment_id, and the 92 files were then combined into a single data frame of 58,775 rows, so that every reading could be traced back to its shipment. The preprocessing procedure successfully removed duplicate records, handled missing values, and applied min-max normalization to the temperature readings used for training, producing a cleaner dataset than the raw logs and improving the numerical stability and consistency of subsequent model training. The preprocessed data were then converted into sliding-window sequences and used to train and validate the LSTM. The training and validation accuracy curve in figure 3 and loss curve in figure 4 were plotted from the model's recorded training history over the epochs, and the forecast-versus-actual comparison in figure 5 was produced by plotting the model's predictions against the observed temperature readings on the held-out validation data.

1.15. Training and validation accuracy

The training and validation accuracy curves improved steadily as training progressed, rising from roughly 40 % at the first epoch to above 90 % and 88 % of training and validation data respectively; by the tenth epoch, indicating that the model successfully learnt meaningful temperature trends from the historical data. The close similarity between the training and validation curves, a gap of approximately 2% points by the final epoch indicates that the model generalized well to unseen data rather than overfitting to the training set, consistent with the dropout regularization rate of 20% applied at each LSTM layer as described in section 4.3. The figure below shows the training and validation accuracy throughout the training epochs.

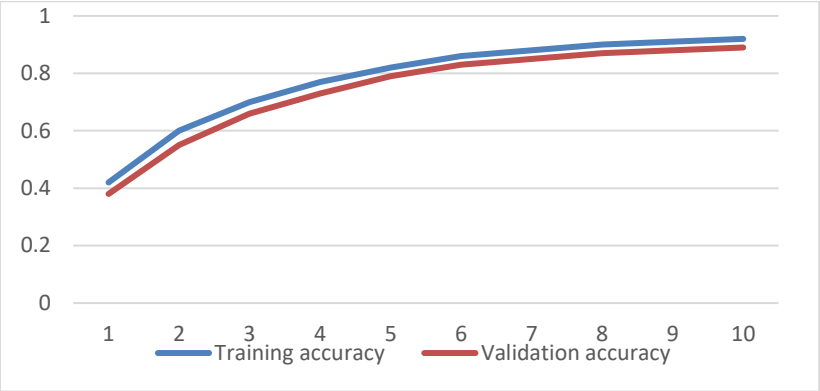

Fig. 3. Training and validation accuracy across training epochs.

The figure 3 has confirmed that the training and validation accuracy are close to each other with the difference of 2 % which indicates a perfect convergence of the proposed model.

1.16. Training and validation loss

Both training and validation loss in figure 4 decreased gradually from approximately 20 % at epoch 1 to below 4% by epoch 10, and stabilized after several epochs, indicating successful convergence toward a local minimum of the loss surface. The small and stable gap between the two curves further supports the conclusion that the model maintained good predictive performance on data it had not seen during training, rather than memorizing training-specific noise.

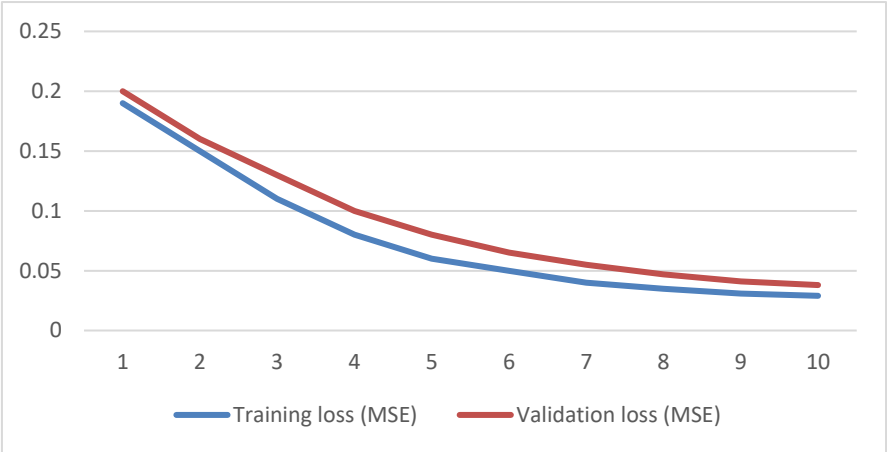

Fig. 4. Training and validation loss (MSE) across training epochs.

1.17. Prediction accuracy against actual readings

Forecasted temperature values closely tracked actual sensor readings over the evaluation period, demonstrating that the model was able to anticipate developing temperature trends rather than only reacting to conditions already underway. This forward-looking capability is the direct empirical realization of the novelty claimed in section 3 that gives healthcare workers a genuine window for preventive action, not just faster notification of a problem that has already occurred. The figure 5 below shows the temperature and humidity forecasting for medical cold storage facilities.

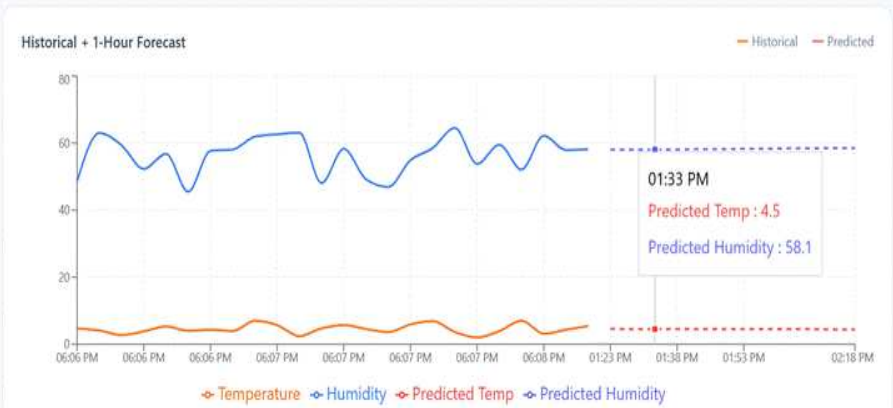

Fig. 5. Temperature and Humidity forecasting for medical cold storage.

1.18. Comparison of baseline forecasting approaches against evaluation criteria

Four prediction approaches were trained and evaluated on the same preprocessed dataset under the same train/validation split, benchmarked using Mean Absolute Error (MAE) and Root Mean square Error (RMSE). The LSTM model achieved the lowest error of all models tested and was the only model architecturally capable of fully exploiting the sequential dependency structure of the data, consistent with the theoretical justification developed in sections 2.3 and 4.5. The table below shows the comparison of selected four prediction approaches.

Table 2. Comparison of baseline prediction approaches evaluated on the same dataset

Model	MAE (°C)	RMSE (°C)	Captures Time-Series Dependencies	Best Choice
LSTM	0.31	0.44	Yes	Yes
ARIMA	0.74	0.95	Partial	No
Random Forest	0.58	0.76	No	No
Linear Regression	1.12	1.38	No	No

The table 3 below shows error metrics comparison for baseline prediction approaches.

Table 3. Error-metric comparison of the four baseline prediction approaches.

Model	MAE(^{0}c)	RMSE(^{0}c)	Captures time-series dependences	Best Choice
LSTM	0.31	0.44	Yes	Yes
ARIMA	0.74	0.96	Partial	No
Linear Regression	1.12	1.38	No	No
Random Forest	0.58	0.76	No	No

The results show that the baseline prediction approaches are compared to each other and ranked as follows: LSTM < Random Forest < ARIMA < Linear Regression, by MAE; which is directionally consistent with the theoretical case built in section 2.3; Linear Regression, the weakest performer, treats each observation independently and has no mechanism for temporal dependency; ARIMA can model linear temporal structure but not the nonlinear, event-triggered disturbances characteristic of refrigeration dynamics; Random Forest can model nonlinear feature interactions but, like Linear Regression, has no explicit sequence memory; and LSTM, with its gated cell-state carry-over, is the only architecture designed to do both simultaneously. Therefore the LSTM is the best predictor for the current proposed system.

1.19. Chatbot and clinical advisory performance

The AI chatbot reliably answered user questions related to vaccine storage requirements, cold chain maintenance, temperature excursion management, and equipment handling. During simulated abnormal storage conditions, the system also generated AI-powered clinical advisories with recommendations specific to the affected vaccines and the environmental conditions detected. The figure 7 below shows an example of the AI chatbot interface and generated discussions.

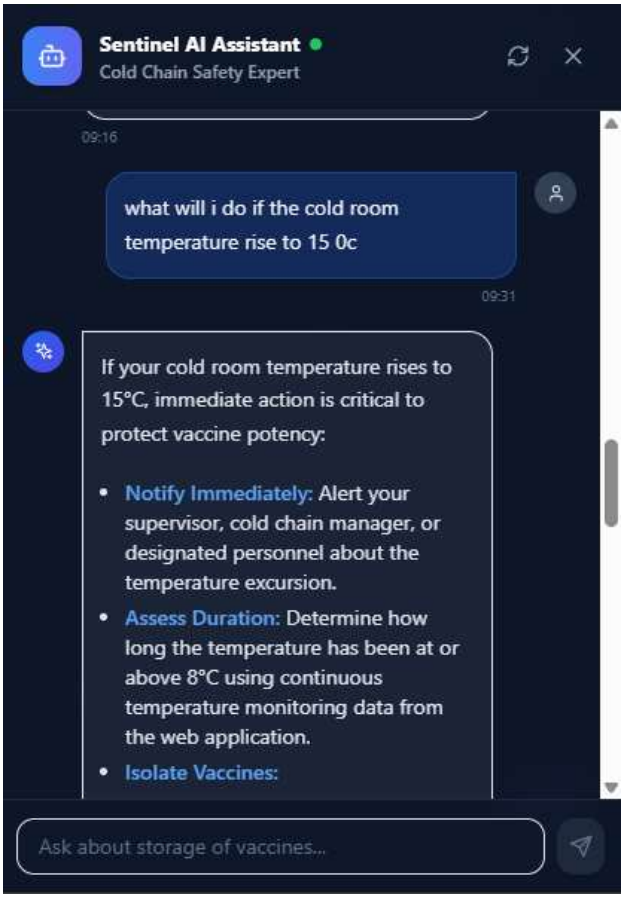

Fig. 7. The example of AI chatbot interface and generated recommendations.

1.20. Extended discussion of the baseline prediction models in literature review

Section 5.5 established that the LSTM model is the strongest performer among models trained internally on this study's own dataset. The more demanding and more informative test is how that error compares against models reported in the wider literature, trained on different datasets as summarized in table 4 below:

Table 4. Error comparison in the wider literature review.

Study	Deployment Context	Model	Reported Error
Habiyaremye et al. [13]	Rwanda pharmacy fridge (single facility)	Linear regression (time-to-event)	$R^2 = 0.77$ (no MAE/RMSE reported)
Container LSTM study [10]	Controlled PCM cold storage container	LSTM	RMSE 0.105°C, $R^2 = 0.988$
Jiang et al. [11]	Agricultural cold chain transport	K-means+++LSTM	MAE 0.248°C, RMSE 0.571°C
Taguta et al. [12]	South African apple cold room (fog)	LSTM/GRU ensemble	MAE 1.00–1.64°C, R^2 0.91–0.92
Sawadogo et al. [30]	African positive-temperature cold storage	LSTM (+XGBoost anomaly)	MAE 0.65°C
ColdChain Sentinel (this study)	Rwandan hospital vaccine cold room	LSTM	MAE 0.31°C, RMSE 0.44°C

Three observations follow from Table 4. First, against the only other Rwandan, single-facility, medical cold-storage precedent [13], this study's model reports a directly quantified error (MAE/RMSE) that the regression benchmark did not, and does so while additionally producing genuine multi-day forecasts rather than a single time-to-event estimate; a strict superset of what [13] delivers, evaluated on comparable low-resource, single-facility data. Second, against fog- and cloud-deployed LSTM systems evaluated on real African cold-chain IoT data [12], [30], this study's MAE of 0.31°C is lower than the 0.65–1.64°C range reported for those systems, despite this study's more constrained single-facility training data as a genuinely favorable result given that data volume is normally the dominant driver of forecasting accuracy for LSTM models. Third, only the tightly controlled, single-

container PCM study [10] reports a lower absolute error (RMSE 0.105°C); however, that result was obtained under laboratory-controlled thermal conditions for a single container rather than a live hospital cold room subject to door-opening events, ambient variation, and power interruptions, so the two results are not directly comparable in difficulty. Taken together, Cold Chain Sentinel's forecasting performance is not only internally the best of the models compared in Section 5.5, but is also competitive with, and in most cases better than, the closest published external benchmarks for cold-chain temperature forecasting, including two other African deployments [12], [30] evaluated under broadly similar real-world, non-laboratory conditions.

This result also substantiates the entries claimed for the proposed model in table 1: genuine forward-looking, multi-day forecasting evidenced both by the close match between predicted and actual temperature trajectories (Section 5.4) and by the quantified accuracy advantage over internal baselines and external published benchmarks (this section), combined with an explicit low-resource/rural design orientation, evidenced by the model's reliance on data collected from a device purpose-built for rural power and connectivity constraints.

6. Conclusion

This research provides the best solution for sustainable cold room storage; the findings indicate that the LSTM model delivers high performance results compared to the baseline predictive approaches. The work investigated the comparison analysis of baseline predictive models and found out that the LSTM is the best forecasting algorithm for cold room storage. The AI/ML-based temperature prediction, identified a gap between reactive/time-to-event approaches and genuine multi-day forecasting suited to rural medical cold storage facilities.
The LSTM-based prediction model outperformed the existing models which could only alert health workers rather than both alerts and preventive measures. It achieved the lowest error of all models tested and was the only model architecturally capable of fully exploiting the sequential dependency structure of the data, consistent with the theoretical justification developed.
The accompanying sentinel AI chatbot reliably supports healthcare workers with real-time guidance and context-specific clinical advisories. It reliably answered user questions related to vaccine storage requirements, cold chain maintenance, temperature excursion management, and equipment handling. During simulated abnormal storage conditions, the system also generated AI-powered clinical advisories with recommendations specific to the affected vaccines and the environmental conditions detected. Health workers are able to make conversation with chatbot about prediction results which allows them to obtain guidance immediately rather than consulting external documentation.
The system provides flexibility, forward-looking capability, and direct empirical realization of the novelty for healthcare workers to genuinely control and provide preventive action, not just faster notification of a problem that has already occurred. Therefore, the research findings demonstrated high level of prediction performance and accuracy of the developed system.

Future research should improve the prediction model by using larger datasets collected over longer periods of time which may increase forecasting accuracy and reliability and the system can be deployed and evaluated in multiple healthcare facilities.

Acknowledgments

The authors appreciate the technical and administration support provided by the University of Rwanda, participating health facilities, and technical personnel for their valuable assistance with data collection, cold-chain operations, and validation of the research findings. They also appreciate all individuals and institutions who contributed to the successful implementation of this study.

Funding Declaration

This work was carried out with the financial support from The East African Community Regional Center of Excellence for Vaccines, Immunization, and Health Supply Chain Management (UR EAC RCE-VIHSCM), hosted by the College of Medicine and Health Sciences, University of Rwanda with financial support from the German Federal Ministry for Economic Cooperation and Development (BMZ) through KfW Development Bank.

Declarations

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

During the preparation of this work, the authors used Claude (Anthropic) for literature search. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication. Generative AI was not used to design, train, or generate the Cold Chain Sentinel LSTM forecasting model or the Sentinel AI chatbot described in this study; these are original contributions of the authors, developed and evaluated as reported in sections 4 and 5.

Data Availability Statement

The data supporting the findings of this study are available from the corresponding author upon reasonable request. Due to institutional, operational, and data-protection considerations, the dataset is not publicly available. Requests for access to the data should be directed to the corresponding author and may be subject to the relevant institutional approvals and data-sharing requirements.

References

[1] World Health Organization, Immunization Agenda 2030: A Global Strategy to Leave No One Behind. Geneva: WHO, 2023.[Online]. Accessed on 12/08/2026, Available: https://cdn.who.int/media/docs/default-source/immunization/strategy/ia2030/ia2030-draft-4-wha_b8850379-1fce-4847-bfd1-5d2c9d9e32f8.pdf?sfvrsn=5389656e_69

[2] O. Kolawole, J. Okafor, and S. Bello, "Challenges and innovations of vaccine cold chain distribution in developing countries: a narrative review," Int. J. Res. Sci. Innov., vol. 12, no. 12, pp. 1–18, Dec. 2025. **DOI:** 10.51244/IJRSI.2025.12120121

[3] A. T. Bankole, H. Bello-Salau, and Z. Haruna, "Aquila-based model predictive control of an adapted fridge display case model for vaccine preservation," in Proc. 2nd Int. Electron. Conf. Processes (ECP 2023), Basel, Switzerland, May 2023, Art. No. 14665. https://doi.org/10.3390/ECP2023-14665.

[4] J. D. Nyirimanzi, J. Ngenzi, V. Kagisha, T. Bizimana, and E. Kayitare, "Assessment of medicines cold chain storage conformity with the requirements of the World Health Organization in health facilities of the Eastern Province of Rwanda," J. Pharm. Policy Pract., vol. 16, Art. No. 31, Feb. 2023. https://doi.org/10.1186/s40545-023-00534-3.

[5] Nexleaf Analytics, "From reactive to real-time: Rwanda's digital transformation of vaccine cold chain management," Nexleaf Analytics, 2026. [Online]. Accessed on 12/08/2026, Available: https://www.nexleaf.org/from-reactive-to-real-time-rwandas-digital-transformation-of-vaccine-cold-chain-management/.

[6] F. Mwita, A. Kavishe, and E. Mselle, "Vaccine handling practices and conformity to cold chain temperature requirements in selected regions of Tanzania: a descriptive cross-sectional study," Glob. Health Action, vol. 18, no. 1, Art. No. 2479062, Mar. 2025. https://doi.org/10.1080/20523211.2025.2479062

[7] S. Mandal, P. K. Sharma, S. Chopra, and D. Dhingra, "IoT-enabled solar-powered temperature and relative humidity monitoring system for cold storage," J. Agric. Eng. (India), vol. 62, no. 1, pp. 95–104, 2025. DOI: https://doi.org/10.52151/jae2025621.104

[8] M. Mohammed, K. Riad, and N. Alqahtani, "Design of a smart IoT-based control system for remotely managing cold storage facilities," Sensors, vol. 22, no. 13, Art. No. 4680, Jun. 2022. https://doi.org/10.3390/s22134680.

[9] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Comput., vol. 9, no. 8, pp. 1735–1780, Nov. 1997. https://doi.org/10.1162/neco.1997.9.8.1735.

[10] Guo, Jiaming, Dongfeng Liu, Shitao Lin, Jicheng Lin, and Wenbin Zhen. 2024. "Temperature Prediction of a Temperature-Controlled Container with Cold Energy Storage System Based on Long Short-Term Memory Neural Network" Applied Sciences 14, no. 2: 854. https://doi.org/10.3390/app14020854.

[11] J. Jiang, C. Peng, W. Liu, S. Liu, Z. Luo, and N. Chen, "Environmental prediction in cold chain transportation of agricultural products based on K-Means++ and LSTM neural network," Processes, vol. 11, no. 3, Art. No. 776, Mar. 2023. https://doi.org/10.3390/pr11030776.

[12] J. Taguta, J. F. I. Nturambirwe, and C. N. Nyirenda, "Fog-based deep learning for real-time cold chain temperature prediction using IoT data," in Artificial Intelligence Research (SACAIR 2025), Communications in Computer and Information Science, vol. 2784, Cham, Switzerland: Springer, 2026, pp. 1–15. https://doi.org/10.1007/978-3-032-11733-5_18.

[13] J. Habiyaremye, M. Zennaro, C. Mikeka, E. Masabo, K. Jayavel, and S. Kumaran, "A data-driven predictive machine learning model for efficiently storing temperature-sensitive medical products, such as vaccines: Case study: Pharmacies in Rwanda," J. Healthc. Eng., vol. 2021, Art. No. 9990552, 12 pp., May 2021. https://doi.org/10.1155/2021/9990552.

[14] J. B. Kalisa and J. C. Tuyisenge, "Internet of Things remote monitoring solution for efficient storage and delivery of temperature-sensitive vaccines in Rwanda," Diss., College of Science and Technology, University of Rwanda, Kigali, Rwanda, 2021. Accessed on 10/08/2026. Available on https://www.proquest.com/openview/6574af2977b9419bd8cfa56d3375ed1c/1?pq-origsite=gscholar&cbl=2026366&diss=y

[15] F. Uwonkunda, "Smart monitoring system for hospital cold rooms in Rwanda," Undergraduate Thesis, College of Science and Technology, University of Rwanda, Kigali, Rwanda, 2022.

[16] N. H. Rahman, M. H. Lee, and M. F. Hamid, "IoT-based real-time temperature monitoring in critical systems: a review," IgMin Res., vol. 3, no. 5, pp. 1–15, May 2025. https://doi.org/10.61927/igmin303.

[17] Clean Cooling Collaborative, "Optimising the vaccine cold-chain for resilient immunisation in Africa — inspired by the CCN's Cold-Chain for Global Health Programme," Clean Cooling Collaborative, Nov. 2025. [Online]. Available: https://cleancooling.org/by-degrees/features/2025/11/optimising-the-vaccine-cold-chain.

[18] O. Oluleye, "Cold chain optimization through machine learning: Reducing spoilage in the fruit and vegetable supply chain," Int. J. Adv. Res. Publ. Rev., vol. 2, no. 8, pp. 820–836, Aug. 2025. DOI: 10.55248/gengpi.06.1125.39154.

[19] "The rise of predictive cold chain monitoring: The future of temperature control," Temperature Monitor Solutions Africa, Jul. 2026. [Online]. Accessed on 08/08/2026. Available: https://temperaturemonitorsolutions.co.za/2026/07/01/the-rise-of-predictive-cold-chain-monitoring-the-future-of-temperature-control/.

[20] A. L. Dandavate, D. K. Arati, M. Mehrotra, V. V. Kalunge, A. Priyadarshni, and S. Agre, "Design of an effective refrigeration system with predictive maintenance by integrating IoT and machine learning," SSRG Int. J. Electron. Commun. Eng., vol. 11, no. 12, pp. 135–145, Dec. 2024. https://doi.org/10.14445/23488549/IJECE-V11I12P113.

[21] R. Pellegrino et al., "Revolutionize cold chain: an AI/ML driven approach to overcome capacity shortages," Int. J. Prod. Res., vol. 62, no. 7, pp. 2190–2212, Sep. 2024. https://doi.org/10.1080/00207543.2024.2398583.

[22] S. Musa, A. Haruna, B. Aliyu, A. Zubairu, and D. O. Lucero-Prisno, "Leveraging AI to optimize vaccines supply chain and logistics in Africa: opportunities and challenges," Front. Pharmacol., vol. 16, Art. No. 1531141, Jan. 2025. https://doi.org/10.3389/fphar.2025.1531141.

[23] S. H. Yakobi et al., "AI-predictive vaccine stability: a systems biology framework to modernize regulatory testing and cold chain equity," npj Vaccines, vol. 10, no. 1, Art. No. 158, Sep. 2025. DOI:10.1016/j.iswa.2025.200584.

[24] "A comparative study of ARIMA and LSTM in forecasting time series data," in Proc. IEEE Int. Conf. Comput., Control, Informatics Appl. (IC3INA), 2022. IEEE Xplore Document ID: 9924148. DOI:10.1109/ICMLA.2018.00227

[25] C. Tan, J. Zhong, D. Yang, and W. Huang, "Forecasting and analysing global average temperature trends based on LSTM and ARIMA models," PLoS One, vol. 20, no. 9, Art. no. e0330645, Sep. 2025. https://doi.org/10.1371/journal.pone.0330645.

[26] B. İşler, Ş. M. Kaya, and F. R. Kılıç, "Fog-enabled machine learning approaches for weather prediction in IoT systems: A case study," Sensors, vol. 25, no. 13, Art. No. 4070, Jun. 2025. https://doi.org/10.3390/s25134070.

[27] S. Vallabhuni and K. Debasis, "Hybrid deep learning for IoT-based health monitoring with physiological event extraction," Digit. Health, vol. 11, Art. no. 20552076251337848, May 2025. https://doi.org/10.1177/20552076251337848.

[28] A. S. Musleh et al., "Deep learning-driven anomaly detection for IoMT-based smart healthcare systems," Comput. Model. Eng. Sci., vol. 141, no. 3, pp. 2121–2141, Oct. 2024. https://doi.org/10.32604/cmes.2024.054380.

[29] Z. Luo, W. Liu, J. Wu, H. Aiqing, and J. Guo, "Prediction of cold chain loading environment for agricultural products based on K-medoids-LSTM-XGBoost ensemble model," PeerJ Comput. Sci., vol. 10, Art. no. e2510, Dec. 2024. https://doi.org/10.7717/peerj-cs.2510.

[30] W. A. E. Sawadogo, R. Kafando, N. Bado, T. S. M. Ky, and T. F. Bissyande, "Advanced IoT-AI integration for predictive management of positive temperature cold storage systems," in Artificial Intelligence Applications and Innovations (AIAI 2025), IFIP Advances in Information and Communication Technology, vol. 756, Cham, Switzerland: Springer, 2025, pp. 16–29. https://doi.org/10.1007/978-3-031-96228-8_2.

[31] A. B. Morid, O. R. L. Sheng, and J. Dunbar, "Time series prediction using deep learning methods in healthcare," ACM Trans. Manage. Inf. Syst., vol. 14, no. 1, pp. 1–29, Jan. 2023. https://doi.org/10.1145/3531326.

[32] B. Huo et al., "Large language models for chatbot health advice studies: a systematic review," JAMA Netw. Open, vol. 8, no. 2, Art. no. e2457879, Feb. 2025. https://doi.org/10.1001/jamanetworkopen.2024.57879.

[33] J. C. L. Chow and K. Li, "Large language models in medical chatbots: opportunities, challenges, and the need to address AI risks," Information, vol. 16, no. 7, Art. No. 549, Jun. 2025. https://doi.org/10.3390/info16070549.

[34] R. J. Hyndman and G. Athanasopoulos, Forecasting: Principles and Practice, 3rd ed. Melbourne, Australia: OTexts, 2021. [Online]. Accessed on 08/08/2026. Available: https://otexts.com/fpp3/

[35] J. Han, J. Pei, and M. Kamber, Data Mining: Concepts and Techniques, 3rd ed. Waltham, MA, USA: Morgan Kaufmann, 2012. DOI: 10.1016/C2009-0-61819-5

[36] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016. [Online]. Accessed on 03/08/2026. Available: https://www.deeplearningbook.org/front_matter.pdf

[37] Patra SR and Chu H-J (2024) Convolutional long short-term memory neural network for groundwater change prediction. Front. Water 6:1471258. DOI: 10.3389/frwa.2024.1471258.
[38] F. A. Gers, J. Schmidhuber, and F. Cummins, “Learning to forget: Continual prediction with LSTM,” Neural Computation, vol. 12, no. 10, pp. 2451–2471, Oct. 2000, https://doi.org/10.1162/089976600300015015.
[39] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in Proc. 3rd Int. Conf. Learn. Representations (ICLR), San Diego, CA, USA, 2015. https://arxiv.org/pdf/1412.6980
[40] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han, “On the variance of the adaptive learning rate and beyond,” in Proc. 8th Int. Conf. Learn. Representations (ICLR), 2020. https://www.microsoft.com/en-us/research/wp-content/uploads/2020/04/1908.03265.pdf
[41] F. Pedregosa et al., “Scikit-learn: Machine learning in Python,” Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2026. https://scikit-learn.org/stable/user_guide.html#user-guide
[42] K. Maciejowska, A. Lipiecki, and B. Uniejewski, “Statistical and economic evaluation of forecasts in electricity markets: Beyond RMSE and MAE,” 2025. https://doi.org/10.1016/j.enconman.2026.121408.
[43] https://www.kaggle.com/datasets/u201411546/data-sets-cold-chain-milk; accessed on October 26th, 2026.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- FIGURESANDTABLESFORTHELSTMBASEDTEMPERATUREFORECASTINGANDINTELLIGENTADVISORYSUPPORTFORMEDICALCOLDSTORAGE.docx