Intelligent system for diabetes management on mobile devices

Victoria Vysotska[1,†], Kirill Smelyakov[2,*,†], Nataliia Sharonova[3,†], Oleksandr Dolhanenko[2,†], Oleksiy Lanovyy[2,†], Vadym Repikhov[2,†]

[1] *Lviv Polytechnic National University, Stepan Bandera Street, 12, Lviv, 79013, Ukraine*
[2] *Kharkiv National University of Radio Electronics, 14 Nauky Ave., Kharkiv, 61166, Ukraine*
[3] *National Technical University "KhPI", Kyrpychova str. 2, Kharkiv, 61002, Ukraine*

Abstract

Effective diabetes management requires continuous monitoring and accurate prediction of blood glucose levels. This research presents an intelligent, mobile-based glucose prediction system that integrates deep learning models, continuous glucose monitoring (CGM) data, and natural language processing (NLP) techniques for automated meal and insulin intake logging. The proposed approach employs Long Short-Term Memory (LSTM) networks to capture temporal dependencies in glucose fluctuations while leveraging large language models (LLMs) to process free-form user inputs. The system aggregates CGM sensor readings, dietary records, and time-based features to enhance prediction accuracy and personalise forecasts. A dedicated mobile application facilitates real-time monitoring and alerts, enabling proactive diabetes management. Experimental evaluation of the system demonstrates its capability to minimise data loss, enhance prediction precision, and improve usability in real-world scenarios. The results indicate a trend of improved accuracy with personalised models, suggesting that integrating AI-driven automation in glucose tracking can significantly benefit diabetes care. Future work will focus on expanding feature integration, refining meal logging capabilities, and conducting clinical validation to ensure broader applicability and regulatory compliance.

Keywords

artificial intelligence, continuous glucose monitoring, deep learning, diabetes management, glucose prediction, information technologies, long short-term memory, machine learning, mobile health, natural language processing, predictive analytics.

1. Introduction

Diabetes management relies on careful monitoring and control of blood glucose levels to prevent dangerous hypoglycemia (low blood sugar) or hyperglycemia (high blood sugar). Continuous Glucose Monitoring (CGM) devices have transformed this process by providing frequent, automatic readings of glucose levels [1]. These CGM readings are typically relayed to smartphone applications, giving users real-time information on their glycemic trends [2].

The real-time data stream enables not only immediate alerts for high or low values, but also opens the door to predicting future glucose levels before critical events occur. Accurate short-term blood glucose prediction is increasingly recognised as a key aspect of diabetes care.

By forecasting where glucose levels are heading in the next 5, 10, or 30 minutes, patients and caregivers can take proactive measures (such as adjusting insulin or consuming carbohydrates) to maintain glucose in a safe range [3].

PhD Workshop on Artificial Intelligence in Computer Science at 9th International Conference on Computational Linguistics and Intelligent Systems (CoLInS-2025), May 15–16, 2025, Kharkiv, Ukraine

* Corresponding author.
† These authors contributed equally.
victoria.a.vysotska@lpnu.ua (V. Vysotska); kyrylo.smelyakov@nure.ua (K. Smelyakov); nvsharonova@ukr.net (N. Sharonova); oleksandr.dolhanenko@nure.ua (O. Dolhanenko); oleksiy.lanovyy@nure.ua (O. Lanovyy); vadym.repikhov@nure.ua (V. Repikhov);
0000-0001-6417-3689 (V. Vysotska); 0000-0001-9938-5489 (K. Smelyakov); 0000-0002-3996-6940 (O. Dolhanenko); 0000-0002-4504-4301 (O. Lanovyy); 0000-0002-1274-4205 (V. Repikhov);

A variety of modelling techniques have been applied to CGM data for glucose forecasting. These range from classical time-series models like ARIMA and exponential smoothing to advanced machine-learning approaches, including recurrent neural networks (LSTM/GRU) and temporal convolutional networks. Modern smartphone hardware, especially with the advancements in CPU, is powerful enough to run many of these predictive algorithms in real-time, meaning personalised prediction models can potentially run on the patient's mobile device [4]. In this article, we review the most popular approaches to glucose data aggregation and forecasting, focusing on providing an automated and intelligent diabetes management solution from the context of modern mobile devices.

The goal of this research is to design and develop a personalised glucose prediction system that leverages deep learning to enhance real-time diabetes management on mobile devices. By integrating historical glucose data, time-based features, and natural language, as well as text and audio processing for automated meal and insulin intake logging, the model aims to capture individual metabolic responses and improve forecasting accuracy. The study focuses on constructing an LSTM-based predictive model, structuring spoken meal logs into standardised records using NLP and LLMs, and optimising real-time adaptability for both on-device and cloud-based inference. This research bridges the gap between AI-driven glucose forecasting and practical diabetes care, making monitoring more precise and accessible — all while running on devices the patients already own.

2. Related works

Continuous Glucose Monitoring (CGM) systems have revolutionised diabetes management by enabling real-time glucose level tracking. These systems, including Dexcom, FreeStyle Libre, and Eversense, are typically integrated with mobile applications such as XDrip and Juggluco, providing both patients and healthcare professionals with detailed insights into glucose trends. These systems track interstitial glucose levels using minimally invasive sensors. The sensors typically record glucose readings every 5 minutes (in some cases, every minute), resulting in 288 (or up to 1440) data points per day for a single user.

Applications such as XDrip and Juggluco enhance data visualisation and storage, enabling both retrospective analysis and real-time decision-making—an essential component of digital diabetes management logs.

CGM time series data reveals critical glucose dynamics, including trends (gradual glucose changes) and anomalies (sudden spikes or drops) [5]. These fluctuations vary among individuals and are influenced by diet, stress, and sensor calibration. Analysing these trends and anomalies helps develop predictive models for anticipating critical glucose changes and improving personalised diabetes management.

2.1. CGM data processing

Modern CGM systems are designed with seamless data transmission to mobile devices via Bluetooth connectivity [6]. This integration enhances real-time patient alerts, facilitates trend monitoring, and enables advanced analytics generation for better diabetes management.

Open-source tools, such as Juggluco, have significantly simplified sensor data aggregation, making CGM data more accessible to a broader range of users and developers [7]. These tools allow users to integrate CGM data from multiple sources into their custom applications or monitoring systems, providing greater flexibility and personalisation in glucose tracking.

However, CGM device manufacturers, such as Freestyle Libre, are continuously strengthening data access controls. With each new device iteration, additional authorisation and security mechanisms are introduced, restricting third-party applications from directly accessing sensor data.

While this restrictive approach is often justified by concerns over security and health data privacy [8], it simultaneously limits opportunities for innovation in independent software development. These barriers make it challenging for developers to create new, AI-powered solutions for glucose prediction and personalised diabetes care without official manufacturer support.

Systems integrated with CGM frequently encounter data gaps, which may arise due to technical failures, environmental factors, or human-related issues. Common causes include:

- Sensor disconnections from the mobile device (some sensors do not support local history caching).
- Battery depletion, leading to data loss during downtime.
- Intermittent signal transmission failures cause incomplete or missing data points.

Additionally, behavioural factors contribute to data inconsistencies:

- Irregular device usage by the user (e.g., removing the sensor periodically).
- Improper sensor calibration is essential before use [9]. Sensors require a warm-up period and initial calibration against a traditional blood glucose meter to ensure accuracy.

These data gaps present significant challenges for predictive modelling, as missing values can lead to reduced model accuracy. Without proper data restoration, glucose prediction models lose reliability and fail to capture real trends.

To ensure data integrity, robust imputation methods are required to effectively recover missing glucose readings. Figure 1 [10] illustrates an example of CGM data imputation using linear interpolation, a common technique for filling data gaps by estimating missing values based on surrounding observations.

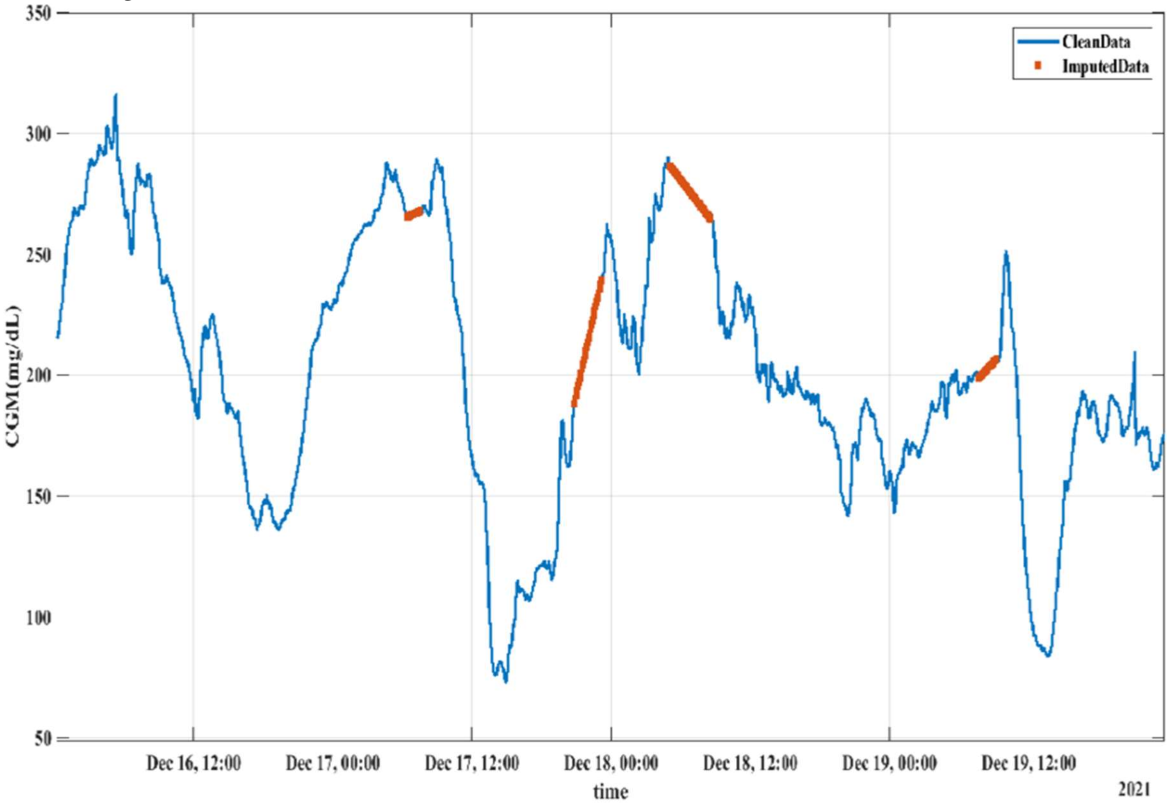

Figure 1: CGM data imputation using Linear Interpolation [10].

The accuracy of CGM sensors depends on both physiological and technical factors. The dynamics of interstitial glucose do not always precisely correspond to blood glucose levels, leading to measurement discrepancies. It occurs because glucose diffusion between blood and interstitial fluid happens with a delay, affecting real-time readings. Additionally, sensor accuracy decreases over time, requiring periodic calibration to ensure correct measurements.

Another major challenge in CGM data analysis is noise [11], which complicates the identification of glucose trends. Fluctuations in glucose levels can result from various influences, including the timing and composition of meals, the impact of physical activity on metabolism, external stress factors, and sleep patterns, which introduce additional variability. These uncontrolled influences make it difficult to extract meaningful trends, necessitating the use of advanced filtering and preprocessing techniques in predictive systems.

2.2. Methods of Glucose Prediction

Several approaches can be used to predict glucose trends. They vary in complexity, accuracy, precision, and scalability, among other factors. Therefore, it is vital to consider and select an appropriate prediction method in order to implement an effective glucose management system.

2.2.1. Autoregressive Integrated Moving Average (ARIMA)

The Autoregressive Integrated Moving Average (ARIMA) model, a traditional statistical approach, is widely used for time series forecasting, especially in identifying short-term glucose level trends. Its versatility allows it to handle both stationary and non-stationary data, making it suitable for analysing glucose dynamics over short periods.

Notably, certain studies [12] have introduced an ARIMA model with adaptive order selection, which enhances the accuracy of blood glucose concentration predictions and improves the detection of hypoglycemia.

This method consists of two key stages. First, the data undergo differentiation to make them stationary, effectively removing trends and stabilising statistical properties. After this transformation, the model forecasts values by combining the autoregressive (AR) component and the Moving Average (MA) component, which utilises past random errors.

The ARIMA model adaptively incorporates past trends and errors, making it highly effective in forecasting glucose levels even in complex and dynamic conditions. However, it faces certain limitations, particularly in handling nonlinear and high-frequency glucose fluctuations [13]. These challenges often require additional preprocessing steps, such as differentiation and seasonal decomposition, to enhance predictive accuracy.

2.2.2. Exponential Smoothing

Exponential smoothing is a statistical method used to smooth discrete time series data, such as blood glucose levels measured at regular intervals. This approach is simple yet effective, capable of adapting to changes in data dynamics while maintaining reasonable accuracy [14]. Its effectiveness lies in the weighted averaging of past observations, where more recent values are assigned greater weight.

Exponential smoothing is particularly useful for continuous glucose monitoring (CGM) systems, where glucose levels are recorded at intervals of 5 to 15 minutes. This method helps filter out noise and irregularities in the data caused by external factors, such as food intake, physical activity, or sensor errors.

Higher values of the parameter α increase the model's sensitivity to recent data changes, whereas lower values contribute to the formation of more stable and smoothed forecasts, emphasising long-term trends. The exponential smoothing method helps to reduce noise in the data while preserving key trends, making it widely applicable for short-term forecasting.

2.2.3. Long Short-Term Memory (LSTM)

To understand the architecture and capabilities of LSTM over other methods, it is essential to introduce some fundamental concepts first. As with any basic neural network, the architecture [15] consists of three main layers:

- Input layer — determines the number of features in the dataset.
- Hidden layers — process data using weighted connections, known as synapses, and activation functions such as sigmoid or tanh.
- Output layer — produces the final prediction while minimising the error between expected and actual values.

The learning process occurs through an iterative optimisation technique called backpropagation [16], which repeatedly adjusts the weights until an optimal accuracy level is achieved.

A Recurrent Neural Network (RNN) is a type of neural network designed explicitly for sequence prediction [17]. In this task, each output depends on the steps taken at the previous time. The hidden layers in an RNN function as a memory, retaining information from earlier steps. It allows the network to identify temporal patterns and trends. However, traditional RNNs face a challenge in maintaining long-term dependencies. It is because of a phenomenon known as vanishing or exploding gradients. During backpropagation, gradients become excessively small or large, which makes learning inefficient.

To overcome this limitation, LSTM (Long Short-Term Memory) and GRU (Gated Recurrent Unit) were developed as advanced types of RNNs. These models introduce gates that regulate the flow of information into and out of the hidden state, allowing the network to learn what to remember and what to forget.

Thus, LSTM is an enhanced version of RNN that is capable of storing long sequences of data. By integrating memory gates, LSTM effectively retains crucial information, leading to more accurate predictions [18].

2.2.4. Temporal Convolutional Networks (TCN)

The Temporal Convolutional Network (TCN) architecture is a powerful tool for glucose level prediction, offering the ability to integrate multiple data sources, such as CGM readings, insulin doses, and carbohydrate intake. The key advantage of TCNs lies in their ability to analyse time series data while capturing both short-term and long-term dependencies [19].

Unlike traditional recurrent architectures, TCNs utilise dilated convolutional layers, which allow the model to cover long temporal intervals without losing computational efficiency. This structure enables TCNs to process multi-dimensional data streams as a single time series, effectively identifying patterns across different time scales.

By integrating contextual data such as physical activity, stress levels, and heart rate, TCNs can be used to build personalised glucose prediction models that adapt to individual physiological characteristics, ultimately enhancing model accuracy [20].

It is important to note that typically, TCN does not replace LSTM but rather complements it, adding layer for processing multi-dimensional input data.

Temporal Convolutional Networks (TCNs) provide a robust and efficient alternative to traditional recurrent neural networks (RNNs). However, it's worth noting that the effectiveness and predictiveness of TCN architecture performance heavily depend on the amount of data available for training. These models may underperform when dealing with smaller datasets.

3. Methods and Materials

A crucial aspect of diabetes technology today is the integration of predictive algorithms with mobile platforms for real-time use. CGM systems like the Dexcom G6/G7, Medtronic Guardian, or Abbott FreeStyle Libre stream glucose readings to smartphones or dedicated receivers at intervals as frequent as every 1–5 minutes. The smartphone acts as a data hub and user interface, aggregating the incoming glucose data and often other relevant inputs (for example, manually-entered meal information or insulin doses). Mobile diabetes apps can thus serve as data aggregators, compiling information from multiple devices into one place for analysis. For instance, platforms like Tidepool or Nightscout allow users to see data from their CGM, insulin pump, blood glucose meter, and manual notes all on a unified timeline [21]. This holistic view is very valuable if one wants to feed multiple data streams into a predictive model. However, many current smartphone apps focus on CGM data alone and use simpler trend analysis to issue alerts (such as rate-of-change-based alarms indicating "falling fast" or "rising fast").

The rise of on-device neural networks means that more sophisticated predictions can now happen locally on the phone. From a technical standpoint, implementing neural networks on mobile devices has become feasible through optimised libraries and frameworks. Tools like TensorFlow Lite and Core ML allow a trained model (e.g. an LSTM or TCN) to be converted into a format that runs efficiently on the limited resources of a phone [22].

3.1. Top-level system architecture

The proposed GluComp Android app acts as the central interface, collecting health data, providing real-time predictions, and generating alerts. Its modular design ensures an intuitive user experience, offline functionality, and advanced data visualisation. The backend infrastructure handles data aggregation, trains personalised neural models, and facilitates secure data sharing while ensuring compliance with privacy regulations.

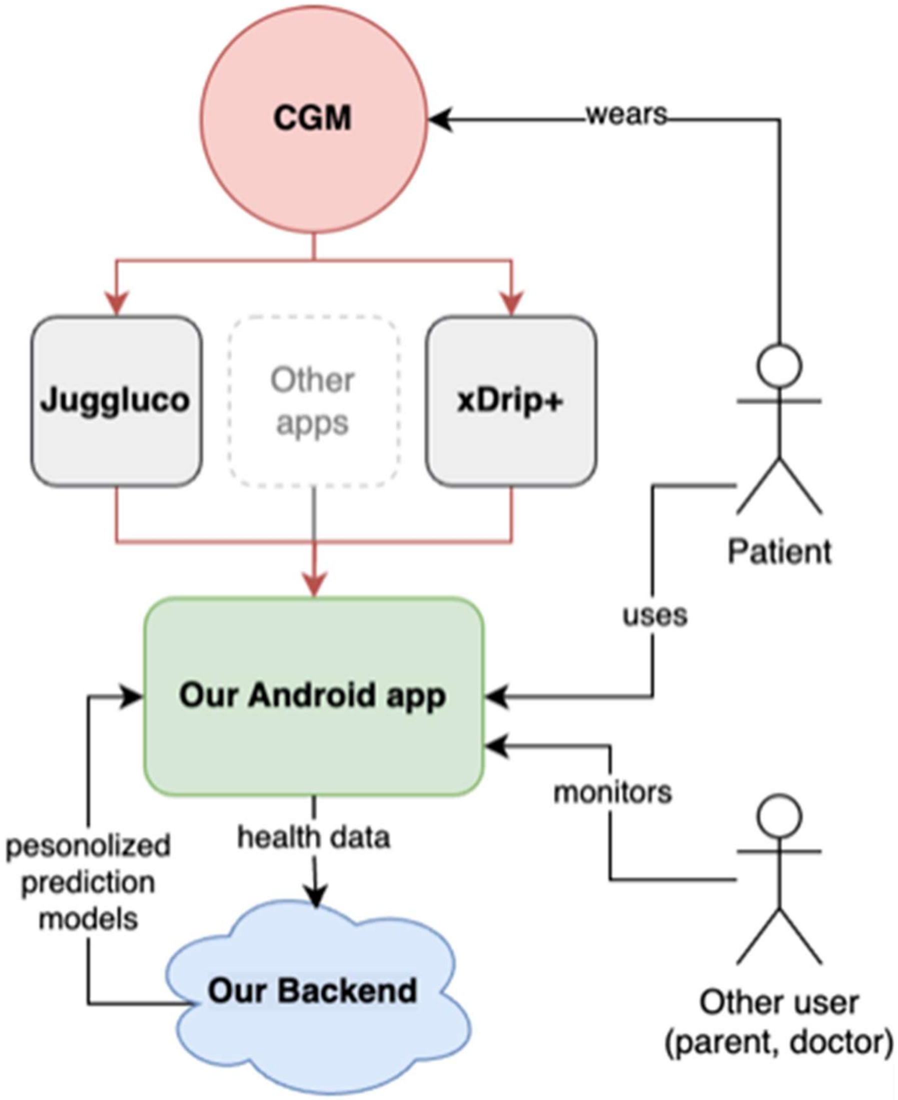

Figure 2: Top-level system architecture.

3.2. Neural network architecture

The proposed model employs deep learning techniques to predict glucose levels based on historical data and contextual time features. By using Long Short-Term Memory (LSTM) networks, the architecture effectively captures temporal dependencies in glucose fluctuations, making it well-suited for personalised predictions.

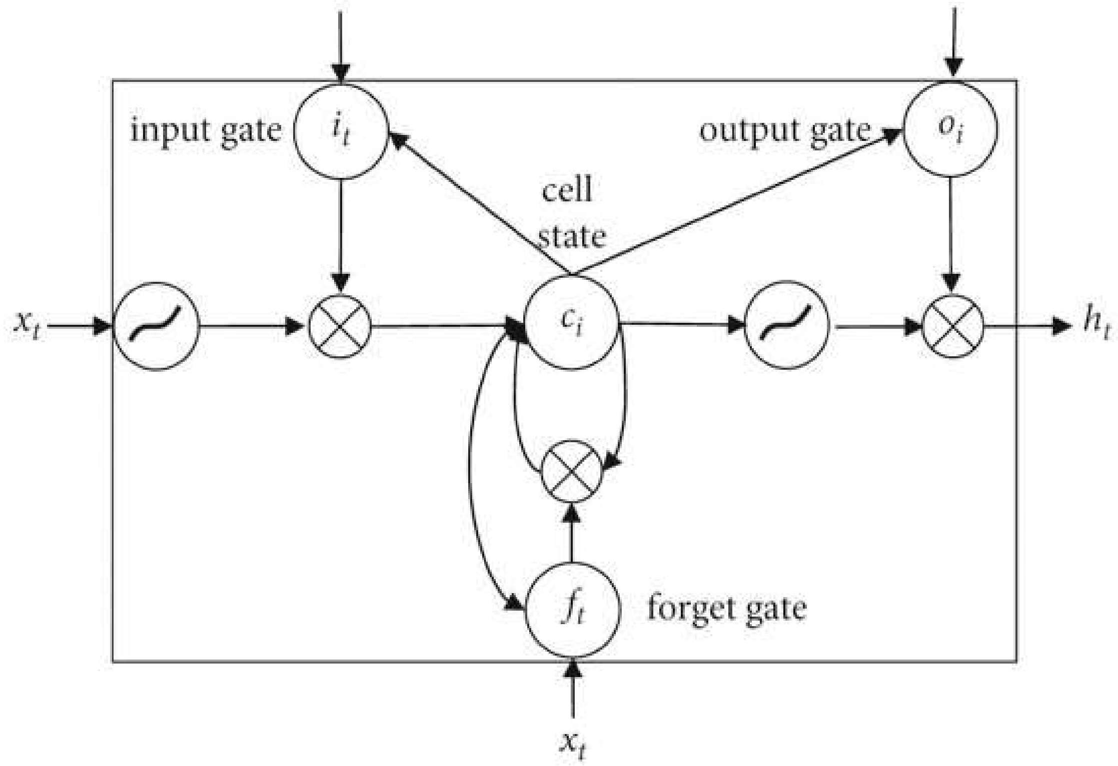

Figure 3: Structure of an LSTM cell [23]

This diagram illustrates the key processes occurring inside an LSTM cell. The input data, along with the previous hidden state and memory cell state, enter the LSTM unit.

The forget gate determines which parts of the previous memory state should be retained or discarded. Functions as a filter, selecting relevant information for the next step.

The input gate processes new information by deciding which input values should be added to the memory cell. The new memory state is formed by combining retained old information and new data added via the input gate.

The output gate uses the updated memory state to compute the final output signal. Determines which part of the memory state should be passed to the hidden state, which is then used for forecasting or forwarded to the next LSTM cell.

We chose LSTM over GRU at this stage, as the model is being created on the backend, where resources are not constrained.

Our model processes two primary inputs. The first input consists of 30 previous glucose values, providing a historical context for trend analysis. The second input incorporates time-based features, such as hour and minute, enabling the model to recognise daily glucose patterns and circadian variations.

The first stage of computation involves an LSTM layer, which learns temporal dependencies from the sequence of past glucose values. Simultaneously, time-based features are encoded using an integer lookup as a multi-hot encoder layer, ensuring that categorical time-related inputs are represented efficiently.

Following these transformations, the outputs from both pathways are concatenated into a unified representation, enabling the model to learn interactions between historical glucose levels and time-dependent variations.

To enhance generalisation and prevent overfitting, a dropout layer is applied before passing the processed features through two hidden layers. These layers refine the learned representations, extracting meaningful patterns that contribute to accurate glucose forecasting.

The final step of the architecture is the output layer, which generates the predicted glucose value. By learning from both sequential trends and time-based influences, the model adapts to individual metabolic patterns, improving the reliability of its forecasts.

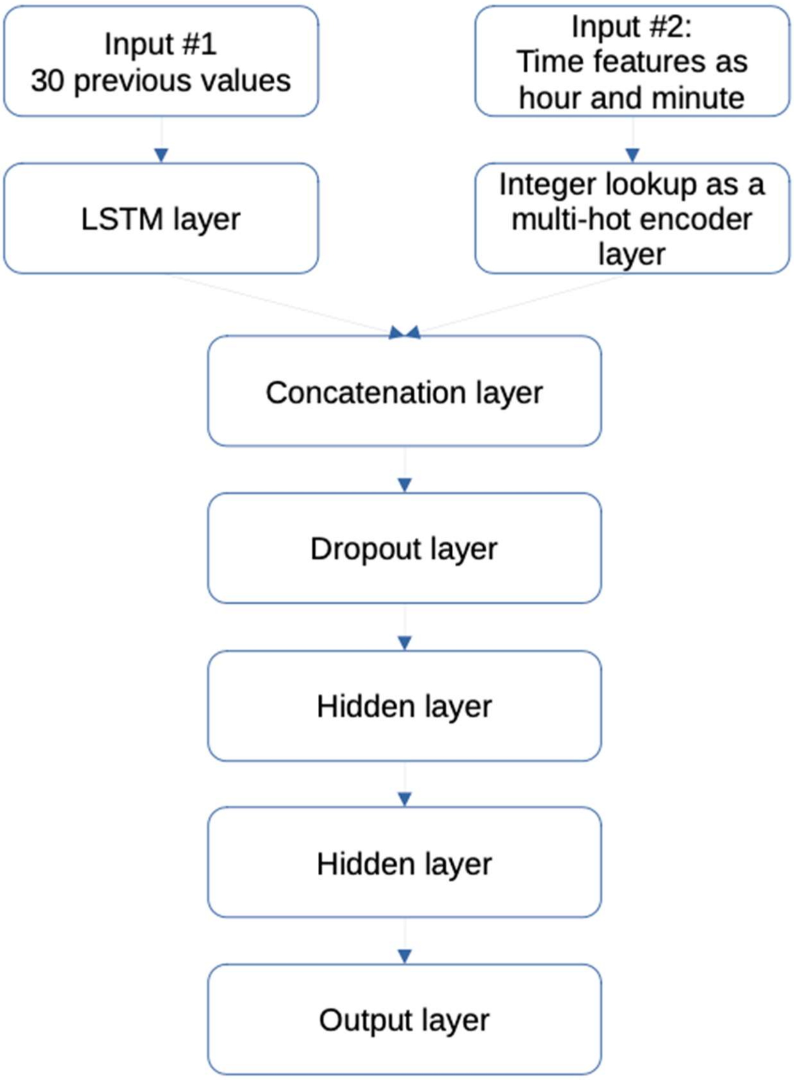

Figure 4: Neural network architecture for predicting glucose values 5 and 10 minutes into the future.

MSE (Mean Squared Error) will be used to evaluate the model's performance.

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2, \tag{1}$$

where n — is the total number of observations;
y_i — is the actual (true) value of the data point;
$\hat{y}_i$ — is the predicted value of the data point;

Additionally, MAE (Mean Absolute Error) is used to further measure the model's performance.

$$MAE = \frac{1}{\mathrm{n}}\sum_{i=1}^{n}|y_i - \hat{y}_i|, \tag{2}$$

where n — the total number of observations;
y_i — the actual (true) value of the data point;
$\hat{y}_i$ — the predicted value of the data point;

These functions will be used after new batches of patient data arrive for further model personalisation.

3.3. Audio and Text Processing for Meal and Insulin Logging

Accurate tracking of food consumption and insulin intake is essential for glucose level prediction and personalised diabetes management. However, manual data entry can be time-consuming and prone to errors. To address this, we propose an automated logging system that leverages natural language processing (NLP) and large language models (LLMs) to process free-form voice recordings [24].

The system can enable users to record an audio log of their meals and insulin intake in a natural and unstructured manner. The audio can then be converted into text on-device and transmitted to the backend for processing, where a combination of NLP techniques and LLM-powered inference extracts and refines the relevant details. In the future, a computer vision approach can be further integrated [25] to simplify the process even more by allowing the user to submit photos of their meal instead of text or audio. It would also mean additional text recognition [26] integrations for understanding insulin packaging labels and doses [27].

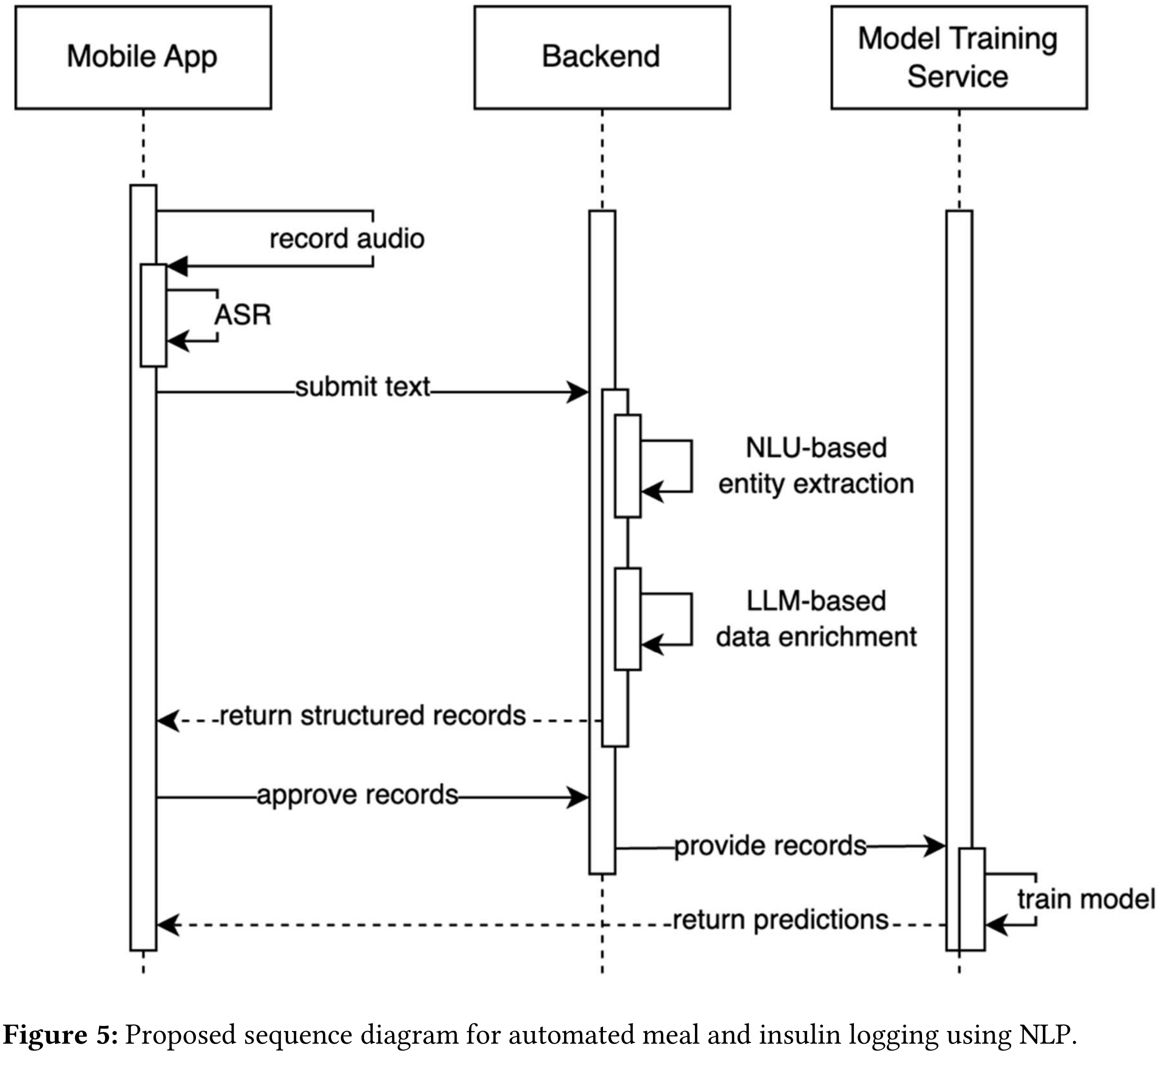

Figure 5: Proposed sequence diagram for automated meal and insulin logging using NLP.

The recorded audio is first transcribed into text using an automatic speech recognition (ASR) system. This step ensures that free-form spoken input is transformed into analysable textual data. The transcribed text undergoes NLP-based entity recognition, extracting key components such as food items, portion sizes, insulin dosage, and meal timing. Context-aware dependency parsing and named entity recognition (NER) help identify structured elements from conversational input. Since users may provide incomplete descriptions (e.g., "I had a bowl of soup"), an LLM-powered inference mechanism is employed to estimate missing nutritional details such as calories, macronutrient composition, and glycemic index. This step ensures that the structured record is both complete and meaningful for future analysis. The processed meal and insulin intake data are then structured into standardised records, categorising each entry with quantified attributes (e.g., meal type, estimated carbohydrate content, insulin dosage). The user can then review and confirm the entries before they are logged into their diabetes management profile. Once verified, the structured nutritional and insulin intake data are fed into the personalised glucose prediction model. By including historical meal and medication patterns, the model continuously adapts to individual metabolic responses, improving long-term predictive accuracy.

As a result, this automated logging system enhances usability, prediction accuracy, and long-term glucose monitoring by reducing manual input efforts while ensuring comprehensive data tracking. Future improvements may include adding projection predicates [28] to support future explainable AI developments, apart from that, personalised LLM fine-tuning to accommodate individual dietary habits and metabolic variations, further optimising prediction outcomes.

4. Experiment

To validate the feasibility of the GluComp system, an experiment was conducted to evaluate the end-to-end functionality of its mobile and backend components [29]. This comprehensive integration test ensured that all subsystems work as intended:

- Sensor connectivity,
- Authentication,
- Data aggregation,
- Cloud synchronisation,
- Prediction precision trends.

The primary objective was to confirm the system's ability to collect glucose data securely, upload it to the cloud, receive machine learning predictions, and present actionable insights to users.

The experiment was split into three separate stages.

4.1. Stage 1: Libre 2 with Juggluco

Stage 1 included 2 weeks of wearing the FreeStyle Libre 2 sensor paired with the Juggluco app v. 8.0.5 transmitting to the GluComp Android application running on Android 15.

The goal was to verify the stability of data transmission from Juggluco to GluComp with a target of < 1% data loss. Key observation points included app performance during night charging periods, where Android process kills occur most often. It was measured by comparing the data entries received by the primary CGM communicator app (Juggluco) and the ones received within GluComp.

Another test point was manual Bluetooth disconnection to simulate physical signal loss with the sensor, where the aim was to verify data recall and re-transmission when the signal appeared again.

4.2. Stage 2: Libre 2 with xDrip+

Stage 2 included 2 weeks of wearing the FreeStyle Libre 2 sensor paired with the xDrip+ app v. dfcbe80-2024.09.17 transmitting to the GluComp Android application running on Android 15. The goal was to verify the stability of data transmission from xDrip+ to GluComp with the same target of < 1% data loss. The observation points from stage 1 apply here as well.

The key difference between Juggluco and xDrip+ in terms of integration is the way they publish data for other applications. Where Juggluco uses a local HTTP server for everything, xDrip+ uses a similar HTTP server for historical queries, and an Android Broadcast system is used to notify about new glucose records in real-time.

It is crucial to test both data transfer technologies (broadcasts and local server), as Android poses certain restrictions on inter-app communication, where, under specific circumstances, the broadcasts may not be delivered.

In a health data aggregator app, data transfer errors should be minimised, as they directly impact health-related decisions.

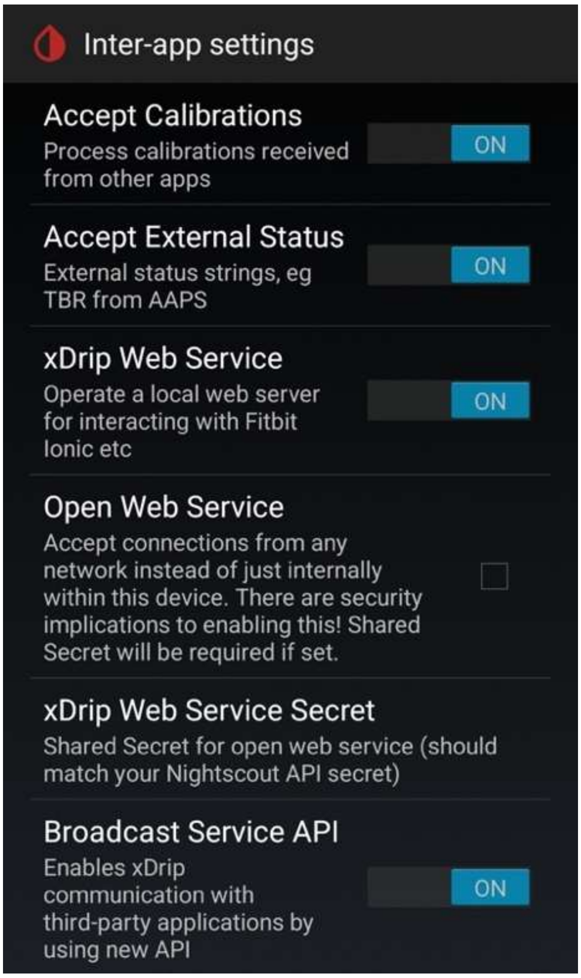

Figure 6: Screenshot from xDrip+ that showcases the settings used for inter-app communication.

Since the method of communication is different, the main goal was to compare the results from this stage with those from stage 1.

4.3. Stage 3: Libre 3 with Juggluco and predictions enabled

Stage 3 included another 2 weeks of wearing the upgraded FreeStyle Libre 3 sensor paired with the Juggluco app v. 8.0.5 transmitting to the GluComp Android application running on Android 15.

The model training service was set to retrain the model every day at 00:00:00 UTC. The goal was to verify that every incoming data point from the CGM resulted in two prediction records being generated (2:1 ratio) with a target of 0% data loss. The additional focus point included calculating the dynamics of the personalised model accuracy and precision. The expectation was that the model accuracy would improve every day during the Stage 3 experiment. The overall accuracy of the predictions was not a concern at this stage, as it requires forming datasets for future experiments.

5. Results

The three stages of the experiment resulted in the following results for the data loss point of interest (see Table 1).

Table 1
Data loss evaluation results

Stage	Sensor	Integration name	Experiment duration (min)	Integration records	Received records	Prediction records
1	Libre 2	Juggluco	20080	6675	6659	N/A
2	Libre 2	xDrip+	20080	3249	3226	N/A
3	Libre 3	Juggluco	20080	19875	19745	36725

The resulting transmission loss is 0.23% for Stage 1, 0.7% for Stage 2 and 0.65% for Stage 3. The resulting prediction loss for Stage 3 is 7%.

The performance trend of the personalised prediction model was compared to that of a general glucose prediction model trained on an open dataset to evaluate the performance trend. The personalised prediction model was additionally trained on 6639 patient records, and the remaining 36 records were used to calculate MSE (Mean Squared Error) and MAE (Mean Absolute Error) [30].

Table 2
MSE and MAE for the general and personalised models

	Mean Squared Error	Mean Absolute Error
General model	681.814	19.065
Personalised model	555.498	16.179

During the experiment, the user interface was also recorded periodically to determine the correctness of the displayed data.

Figure 7 shows the integration screen, where the user can enable the available glucose integration. As seen on the XDrip+ integration section, the "Last update" reveals an issue with the transmission, which was later identified to be the battery optimisation restriction blocking the GluComp service from collecting the CGM data.

The result of the main screen displaying the current and predicted glucose levels can be seen in Figure 8. This screen contains the current glucose value, as well as the historical chart with predictions. The first part of the image shows the Juggluco data source selected, and the second part demonstrates the CGM emulator, resulting in a steady stream of data and predictions. The images reveal the ability to support dark and light themes, as well as different screen sizes.

The application was seen to momentarily adapt to changes in data sources — switching between Juggluco and xDrip+ did not produce additional glucose record gaps, apart from the ones where the sensor was physically disconnected from the device and did not manage to recover upon reconnection.

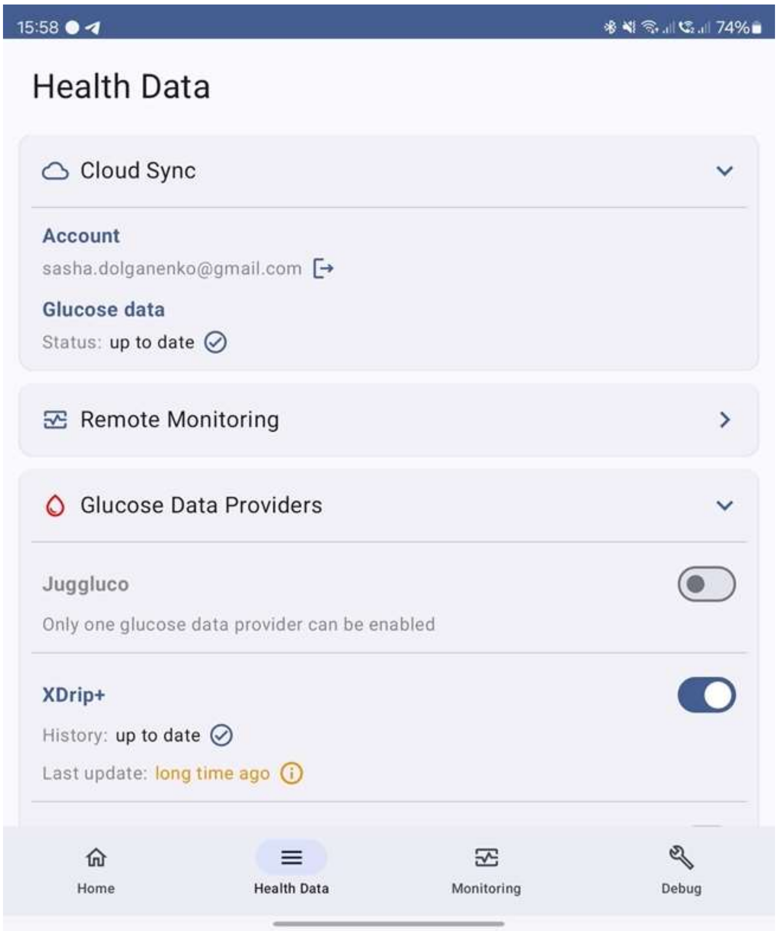

Figure 7: An example of the integration malfunction being displayed due to a connection loss issue.

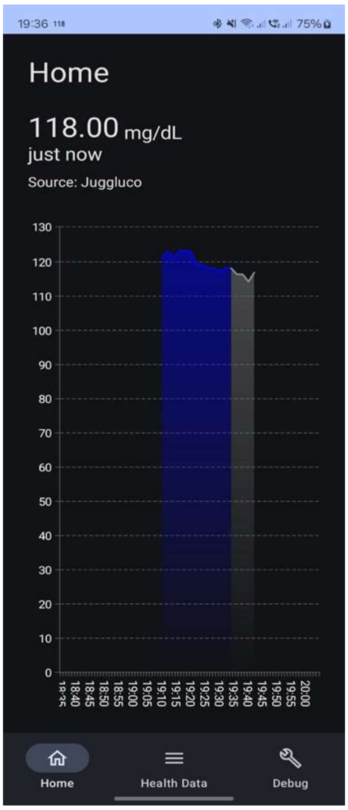

Figure 8: Displaying the current and predicted glucose levels in different UI modes.

6. Discussions

The GluComp Android app integrates with CGM sensors, securely collects and stores glucose data, and provides real-time predictions. The backend processes data, generates predictions and ensures data security and compliance.

The experiment's results showcased the system's full functionality as a glucose monitoring and prediction platform. The successful integration of mobile and backend components validates its potential for real-world deployment and further improvements.

As seen from the results, although integrating both Juggluco and xDrip+ has resulted in data loss, it is kept below 1%. In order to fully mitigate data loss, additional measures need to be taken, such as bypassing battery optimisation settings for GluComp on Android 14+. Moreover, different smartphone manufacturers can set different constraints that can further affect inter-app communication, process-killing policies, and, therefore, data loss. These anomalies should be investigated further.

As for the prediction mechanism, we see quite a stable record generation process, with an output of ~1.86:1 (generated records for every incoming CGM record). The data loss here (compared to the expected 2:1 ratio) is explained by one instance of a failure to export a TensorFlow model to TensorFlow Lite. The resulting exported model produced a day of failures during inference due to a false output structure. Open tickets on TesorFlow GitHub further support this conclusion regarding the incorrect export behaviour. The workaround we will be employing is additional structure verification after export to ensure the resulting .tflite model is correct.

We also see a prediction performance improvement trend when comparing a general model to a personalised model. We anticipate much higher prediction accuracy with the availability of additional data sets and after incorporating extra features in our neural network architecture. Therefore, the next phase involves conducting a closed beta test with real patients to further evaluate the accuracy of the personalised prediction models and introduce additional model features, such as sleep patterns, insulin, and food intake (using automated NLP logging).

7. Conclusion

Neural network-powered glucose prediction on mobile devices represents a convergence of biomedical engineering and personal computing that has great promise for improving diabetes care. Traditional time-series methods like ARIMA and exponential smoothing laid the groundwork for understanding glucose dynamics and are still helpful for quick baseline predictions. Still, they struggle with the complex, nonlinear nature of human metabolism. Advanced models such as LSTMs and GRUs bring memory and learning capabilities that can adapt to individual patient patterns. At the same time, TCNs and other architectures leverage convolutional approaches to model long-term dependencies efficiently. With the advent of robust mobile hardware and software toolkits, these algorithms can be deployed in smartphones or wearables, delivering real-time forecasts to users anytime and anywhere. It enables a shift toward proactive diabetes management – instead of just reacting to current glucose values, patients can get a glimpse into the future and act to prevent excursions before they happen.

We designed, built, and tested our health data aggregator solution, where we focused on minimising data transmission loss from the CGM to our application. Although at this stage, our neural network architecture does not yet support additional features (such as heart rate, diet, or medication timing), we have achieved a positive improvement in prediction accuracy. Integrating natural language processing, computer vision, and LLMs into the system will allow us to achieve seamless insulin and meal intake logging to further improve the model accuracy in the future.

Lastly, regulatory approval and rigorous clinical validation will be essential to ensure both patients and clinicians trust these tools. Current limitations include the need for larger, more diverse patient datasets and long-term studies to confirm real-world effectiveness and safety. Despite these challenges, the trajectory is clear: intelligent mobile systems empowered by neural networks and driven by advances in mobile CPU and GPU technologies are becoming an integral part of diabetes management. These systems promise to help users stay one step ahead of their glucose levels in a way that is both convenient and increasingly effective.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT-4o and Grammarly to check grammar and spelling and perform peer-reviewed simulations. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] K. T. Putra, I. Surahmat, A. N. Nazilah Chamim, M. Z. Ramadhan, D. Wicaksana and R. A. Dhea Namyra Alissa, "Continuous Glucose Monitoring: A Non-Invasive Approach for Improved Daily Healthcare," 2023 3rd International Conference on Electronic and Electrical Engineering and Intelligent System (ICE3IS), Yogyakarta, Indonesia, 2023, pp. 395-400. doi:10.1109/ICE3IS59323.2023.10335328.

[2] S. Gopinath, U. S, S. N. Jathin P, P. V. J and S. Palaniswamy, "Glycemic Index Based Food Recommendation System Using Deep Learning," 2024 1st International Conference on Communications and Computer Science (InCCCS), Bangalore, India, 2024, pp. 1-6. doi:10.1109/InCCCS60947.2024.10593438.

[3] S. -M. Lee, D. -Y. Kim and J. Woo, "Glucose Transformer: Forecasting Glucose Level and Events of Hyperglycemia and Hypoglycemia," in IEEE Journal of Biomedical and Health Informatics, vol. 27, no. 3, pp. 1600-1611, March 2023. doi:10.1109/JBHI.2023.3236822.

[4] A. Das, Y. D. Kwon, J. Chauhan and C. Mascolo, "Enabling On-Device Smartphone GPU based Training: Lessons Learned," 2022 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops), Pisa, Italy, 2022, pp. 533-538. doi:10.1109/PerComWorkshops53856.2022.9767442.

[5] N. Baysal, F. Cameron, B. A. Buckingham, D. M. Wilson and B. W. Bequette, "Detecting sensor and insulin infusion set anomalies in an artificial pancreas," 2013 American Control Conference, Washington, DC, USA, 2013, pp. 2929-2933. doi:10.1109/ACC.2013.6580279.

[6] C. Sreenivas and S. Laha, "Compact Continuous Non-Invasive Blood Glucose Monitoring using Bluetooth," 2019 IEEE Biomedical Circuits and Systems Conference (BioCAS), Nara, Japan, 2019, pp. 1-4. doi:10.1109/BIOCAS.2019.8918744.

[7] D. Margan and S. Čandrlić, "The success of open source software: A review," 2015 38th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO), Opatija, Croatia, 2015, pp. 1463-1468. doi:10.1109/MIPRO.2015.7160503.

[8] S. S. Mahadik et al., "Digital Privacy in Healthcare: State-of-the-Art and Future Vision," in IEEE Access, vol. 12, pp. 84273-84291, 2024. doi:10.1109/ACCESS.2024.3410035.

[9] C. Cenerini, A. Sabatini, L. Vollero and D. Pau, "Optimising Glucose Sensor Calibration With Lightweight Neural Networks: A Comparative Study," in IEEE Sensors Letters, vol. 8, no. 9, pp. 1-4, Sept. 2024, Art no. 6010904. doi:10.1109/LSENS.2024.3436630.

[10] Butt, Hatim & Khosa, Ikramullah & Iftikhar, Muhammad. (2023). Feature Transformation for Efficient Blood Glucose Prediction in Type 1 Diabetes Mellitus Patients. Diagnostics. 13. 340. doi:10.3390/diagnostics13030340.

[11] N. Bidarahalli, M. Rajesh, V. Velappan and P. T. Krishnan, "Noise Reduction in Continuous Blood Glucose Sensor using Physiology based Kalman Filter for Artificial Pancreatic System," 2018 3rd International Conference on Circuits, Control, Communication and Computing (I4C), Bangalore, India, 2018, pp. 1-4. doi:10.1109/CIMCA.2018.8739339.

[12] J. Yang, L. Li, Y. Shi and X. Xie, "An ARIMA Model With Adaptive Orders for Predicting Blood Glucose Concentrations and Hypoglycemia," in IEEE Journal of Biomedical and Health Informatics, vol. 23, no. 3, pp. 1251-1260, May 2019. doi:10.1109/JBHI.2018.2840690.
[13] Siami-Namini, S., Tavakoli, N., & Siami Namin, A. (2018). A Comparison of ARIMA and LSTM in Forecasting Time Series. 2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA). doi:10.1109/icmla.2018.00227.
[14] M. A. F. Quioc, S. C. Ambat, A. C. Lagman, R. F. Ramos and R. R. Maaliw, "Analysis of Exponential Smoothing Forecasting Model of Medical Cases for Resource Allocation Recommender System," 2022 10th International Conference on Information and Education Technology (ICIET), Matsue, Japan, 2022, pp. 390-397. doi:10.1109/ICIET55102.2022.9778987.
[15] M. A. M. Sadeeq and A. M. Abdulazeez, "Neural Networks Architectures Design, and Applications: A Review," 2020 International Conference on Advanced Science and Engineering (ICOASE), Duhok, Iraq, 2020, pp. 199-204. doi:10.1109/ICOASE51841.2020.9436582.
[16] M. Dampfhoffer, T. Mesquida, A. Valentian and L. Anghel, "Backpropagation-Based Learning Techniques for Deep Spiking Neural Networks: A Survey," in IEEE Transactions on Neural Networks and Learning Systems, vol. 35, no. 9, pp. 11906-11921, Sept. 2024. doi:10.1109/TNNLS.2023.3263008.
[17] Y. Chen and J. Li, "Recurrent Neural Networks algorithms and applications," 2021 2nd International Conference on Big Data & Artificial Intelligence & Software Engineering (ICBASE), Zhuhai, China, 2021, pp. 38-43. doi:10.1109/ICBASE53849.2021.00015.
[18] X. Wen and W. Li, "Time Series Prediction Based on LSTM-Attention-LSTM Model," in IEEE Access, vol. 11, pp. 48322-48331, 2023. doi:10.1109/ACCESS.2023.3276628.
[19] S. Chauhan, Sehaj, S. Kumar, Siddharth, D. Panwar and S. Chopra, "Temporal Convolutional Network and its Application in Various Sectors," 2023 3rd Asian Conference on Innovation in Technology (ASIANCON), Ravet IN, India, 2023, pp. 1-7. doi:10.1109/ASIANCON58793.2023.10270242.
[20] Aashima, S. Bhargav, S. Kaushik and V. Dutt, "Temporal Convolutional Networks Involving Multi-Patient Approach for Blood Glucose Level Predictions," 2021 International Conference on Computational Performance Evaluation (ComPE), Shillong, India, 2021, pp. 288-294. doi:10.1109/ComPE53109.2021.9752461.
[21] Arduser, Lora. (2017). Impatient Patients: A DIY Usability Approach in Diabetes Wearable Technologies. Communication Design Quarterly. 5. 31-39. doi:10.1145/3188387.3188390.
[22] M. C. Chirodea, O. C. Novac, C. M. Novac, N. Bizon, M. Oproescu and C. E. Gordan, "Comparison of Tensorflow and PyTorch in Convolutional Neural Network - based Applications," 2021 13th International Conference on Electronics, Computers and Artificial Intelligence (ECAI), Pitesti, Romania, 2021, pp. 1-6. doi:10.1109/ECAI52376.2021.9515098.
[23] Liu, Bingru & Zhao, Qing. (2022). Implementation of Financial Audited Robot Question and Answer Technology of Feature Processing and Improved Bi-LSTM. Scientific Programming. 2022. 1-10. doi:10.1155/2022/7213395.
[24] K. Smelyakov, D. Karachevtsev, D. Kulemza, Y. Samoilenko, O. Patlan and A. Chupryna, "Effectiveness of Preprocessing Algorithms for Natural Language Processing Applications," 2020 IEEE International Conference on Problems of Infocommunications. Science and Technology (PIC S&T), Kharkiv, Ukraine, 2020, pp. 187-191. doi:10.1109/PICST51311.2020.9467919.
[25] Byzkrovnyi, O., Savulioniene, L., Smelyakov, K., Sakalys, P., Chupryna, A. Comparison of Potential Road Accident Detection Algorithms for Modern Machine Vision System, Vide. Tehnologija. Resursi - Environment, Technology, Resources, 2023, 3, pp. 50–55. doi:10.17770/etr2023vol3.7299.
[26] Smelyakov, K., Chupryna, A., Darahan, D., Midina, S. Effectiveness of modern text recognition solutions and tools for common data sources, CEUR Workshop Proceedings, 2021, 2870, pp. 154–165. URL: https://ceur-ws.org/Vol-2870/.

[27] D. Uhryn, V. Vysotska, D. Zadorozhna, M. Spodaryk, K. Hazdiuk, Z. Hu, Intelligent Application for Predicting Diabetes Spread Risk in the World Based on Machine Learning. International Journal of Intelligent Systems and Applications(IJISA), Vol.17, No.3, pp.90-144, 2025. DOI:10.5815/ijisa.2025.03.06

[28] G. Proniuk, N., Geseleva, I. Kyrychenko, G. Tereshchenko, Spatial interpretation of the notion of relation and its application in the system of artificial intelligence, CEUR-WS, 2019, v. 2362, pp. 266-276. ISSN 16130073.

[29] C. Augusto, "Efficient test execution in End to End testing : Resource optimisation in End to End testing through a smart resource characterisation and orchestration," 2020 IEEE/ACM 42nd International Conference on Software Engineering: Companion Proceedings (ICSE-Companion), Seoul, Korea (South), 2020, pp. 152-154.

[30] H. D. P. D. Carvalho, W. L. Soares, W. B. Santos and R. Fagundes, "A Comparison Study About Parameter Optimisation Using Swarm Algorithms," in IEEE Access, vol. 10, pp. 55488-55498, 2022. doi:10.1109/ACCESS.2022.3175202.