Real-Time Blood Glucose Prediction Using SARIMAX and Edge Computing

Wasif Arshad
Department of Electrical Engineering, The Superior University Lahore
arshadwasif2022@gmail.com

Murtaza Anjum
Department of Electrical Engineering, The Superior University Lahore

Mustafa Shakir
Department of Electrical Engineering, The Superior University Lahore
Corresponding Author: * Wasif Arshad arshadwasif2022@gmail.com

ABSTRACT

The global surge in blood glucose regulation disorders has increased demand for smarter, predictive healthcare systems. While Continuous Glucose Monitoring (CGM) devices have improved real-time tracking of glucose levels, their reactive nature limits proactive management. Anticipating glycemic fluctuations remains a significant challenge in minimizing health risks for individuals affected by these fluctuations. This study introduces a transparent and lightweight forecasting framework using the SARIMAX (Seasonal AutoRegressive Integrated Moving Average with eXogenous variables) model, implemented on a Raspberry Pi 5 for real-time execution in low-resource settings. The system employs synthetically generated CGM data over a 3-month window, simulating realistic glucose behavior, including post-meal spikes and circadian trends. The SARIMAX model was trained on this dataset and configured with daily seasonality parameters to forecast glucose levels on an hourly basis over a 7-day prediction horizon. Forecasted values were clipped to remain within physiological boundaries, and uncertainty bands were constructed using an 80% confidence interval. Performance evaluation metrics such as Mean Absolute Error (3.36%), Root Mean Square Error (4.30%), and average deviation (~4.83 mg/dL) confirm the model's high predictive accuracy. Visual overlays of actual and predicted glucose curves further demonstrated close alignment in trend and amplitude, with the majority of readings remaining within clinically acceptable zones. These results validate the model's robustness for short-term forecasting without the need for complex machine learning architectures. The presented framework proves that SARIMAX offers a compelling balance between interpretability, accuracy, and computational efficiency. Its successful deployment on an edge-computing device like Raspberry Pi confirms its viability for real-time applications in patient-centric monitoring systems. This work provides a foundation for future integration of lightweight predictive analytics into mobile healthcare infrastructure and early-warning platforms for glycemic risk mitigation.

Keywords: Glucose Forecasting, SARIMAX Model, Continuous Glucose Monitoring, Raspberry Pi, Real-Time Healthcare, Glycemic Risk Prediction

ACRONYMS

ACF:	Autocorrelation Function
BG:	Blood Glucose
CGM:	Continuous Glucose Monitoring
CPU:	Central Processing Unit
CSV:	Comma-Separated Values
FDA:	Food and Drug Administration
IoT:	Internet Of Things
LCD:	Liquid Crystal Display
MAE:	Mean Absolute Error
PACF:	Partial Autocorrelation Function
PH:	Prediction Horizon
RMSE:	Root Mean Square Error
SARIMA:	Seasonal AutoRegressive Integrated Moving Average
SARIMAX:	SARIMA with eXogenous Variables
SMBG:	Self-Monitoring of Blood Glucose

INTRODUCTION

Diabetes mellitus is a rapidly escalating global health crisis, with more than 537 million people affected worldwide in 2021 and projections indicating a rise to over 783 million by 2045 [1]. This chronic condition is marked by the body's inability to regulate blood glucose (BG) levels, leading to persistent hyperglycemia or dangerous episodes of hypoglycemia if not managed properly [2].

Modern diabetes management heavily relies on Continuous Glucose Monitoring (CGM) systems, which offer near real-time data to track glycemic fluctuations [3]. However, while CGM devices are valuable for observation, they are primarily reactive and lack forecasting capabilities that could allow for preemptive interventions. Therefore, predictive models that estimate future glucose levels are becoming essential to improving patient safety and reducing long-term complications [4].

In recent years, researchers have explored various data-driven forecasting approaches, from statistical methods to advanced machine learning techniques. Among these, deep learning models like Long Short-Term Memory (LSTM) and ensemble methods such as XGBoost have shown strong performance, especially when trained on population-level CGM datasets [5]. However, these methods often require large, high-quality datasets, struggle with interpretability, and are computationally expensive, making them less feasible for deployment in low-resource or real-time settings [6].

As an alternative, Seasonal Auto Regressive Integrated Moving Average with eXogenous variables (SARIMAX) models have gained attention due to their simplicity, transparency, and capacity to model seasonality and exogenous influences such as meal timing or insulin intake [7]. When applied to glucose forecasting tasks, SARIMAX has proven capable of handling cyclical patterns inherent in-patient behavior, especially when reinforced by structured preprocessing techniques such as data clustering and time-segment alignment [8].

This study builds on these developments by evaluating the performance of SARIMAX in forecasting BG levels using mock CGM data that simulates realistic patient conditions. The proposed approach emphasizes model interpretability and computational efficiency while achieving high accuracy. The forecasting results achieved a Mean Absolute Error (MAE) of 3.36%, Root Mean Square Error

(RMSE) of 4.30%, and an average deviation of approximately 4.83 mg/dL, reflecting strong predictive capability suitable for real-time applications.

LITERATURE REVIEW

The use of Continuous Glucose Monitoring (CGM) systems has already made a significant contribution to the way in which diabetic people take care of their health. These systems offer regular readings on the level of blood glucose (BG), and thus prevent risky falls or surges. Nevertheless, CGM devices primarily present present or historical measurements and do not indicate about potential changes in the future. In order to enhance safety, investigators have found various ways of predicting future variations of blood sugar level even before they arise [9].

In some of them, one of the most trustworthy (although this is not accurate) ways is utilizing classic statistical models such as SARIMAX. This model may investigate the common cycles of glucose, including the hourly bodies, and may also take into account the external factors, such as meal schedules or insulin activity [10]. As an illustration of such complexities, one method divides the daily data into chunks by meals or by sleep, so it is simplified to train and test the prediction model based on repeatable data patterns [11]. The ability to predict accurately even when blood sugar behavior varies day to day has also been revealed in other published studies using simulated patient data and similar to the UVA/Padova simulator [12].

SARIMAX models are simpler to interpret and easy to execute as compared to complicated computer models that require extensive data and tend to be a black box. They may also consider the additional information, e.g. the time of a meal intake, in order to enhance the accuracy without overloading the system [13].

It has been demonstrated that even simple models such as ARIMA are generating good results in certain situations. ARIMA has, in fact, been found to be effective in predicting the level of glucose and cholesterol when it is applied in the controlled applications [14]. Differently, another study managed to obtain similar glucose accuracy with SARIMAX models at both short- and mid-term, albeit with fewer input requirements and using fewer computing power compared to more sophisticated models, such as LSTM [15].

These models are frequently checked by simulation tools e.g UVA/Padova simulator prior to their implementation in real environments. Using them researchers are able to test the effectiveness of a given model to address different conditions without committing to using actual patient information [16].

Even though the new forecasting systems are being developed, which integrate artificial intelligence (AI) and historic methods, the experts continue to use models such as SARIMAX. They tend to be more explanatory, simple to describe and work quite well when properly designed. This is why they can be a good choice in the real life healthcare scenarios where simplicity, accuracy before trust is valued most.

In the clinical sphere, medical teams and physicians are more likely to choose the tools which are easy to explain and describe to the patients. SARIMAX models assist in that by showing in a clear manner the influence of regular patterns and daily occurrences on blood sugar and it enhances safety as well as decision-making.

Last but not least, SARIMAX is also appropriate to small devices (low-powered), e.g. Raspberry Pi. It can be implemented in devices such as portable health devices giving real-time updates, which is why it will work well on continuous home monitoring or wearable technologies since it does not require any heavy computing resources.

TABLE 1: Comparison of SARIMA, LSTM, and SARIMAX for Glucose Forecasting

Feature / Model	SARIMA (Seasonal ARIMA)	LSTM (Long Short-Term Memory)	SARIMAX (SARIMA with Exogenous Factors)
Exogenous Factors (X)	No (cannot directly incorporate external variables like meal times, insulin)	Yes (can incorporate, but requires careful data pre-processing and model complexity)	Yes (explicitly and directly models the impact of external variables)
Seasonality Handling	Yes (explicitly models seasonal patterns)	Yes (can learn, but implicitly and requires sufficient data)	Yes (explicitly models seasonal patterns)
Interpretability	Moderate to High (model coefficients can be analyzed)	Low (black-box approach, harder to understand how predictions are made)	Moderate to High (coefficients can be analyzed, exogenous effects are clear)
Non-linearity	Limited (primarily linear relationships)	High (excels at learning complex non-linear relationships)	Limited (primarily linear, unless transformed data)
Data Requirement	Moderate (performs well with less data than LSTMs)	High (requires large datasets for optimal performance)	Moderate (can perform well with moderate data, especially with strong exogenous factors)
Computational Cost	Low to Moderate	High (especially for training)	Low to Moderate
Suitability for Blood Glucose Forecasting	Good for patterns, but misses external influences (e.g., meals)	Can be very accurate with large datasets, but less interpretable for clinical insight	Excellent (explicitly models daily patterns AND crucial external drivers like meals)

Table 1 provides a comparison of three approaches to blood glucose predictions, SARIMA, LSTM, and SARIMAX, according to their major characteristics, such as accuracy, usability, and potential to incorporate external dependencies, such as meals. SARIMA can be aptly used to identify the patterns daily but fails to accommodate the outside factors. The LSTM will be able to perform on more complicated patterns, however requires a significant amount of data, and is less understandable. SARIMAX is a middle ground in the sense that it contains both the daily trend as well as external influences but will be relatively simple to use and interpret. It can also be applied on miniature devices such as Raspberry Pi, which makes it an efficient choice as far as health monitoring in-real-time is concerned.

METHODOLOGY

This section contains the description of the system design, preparation of the input data, the model architecture, and the environment in which the implementation of the proposed system of blood glucose forecasting using SARIMAX on Raspberry Pi 5 was to be conducted as well as the methodology of testing the proposed system of the forecast of blood glucose level on Raspberry Pi 5 using SARIMAX.
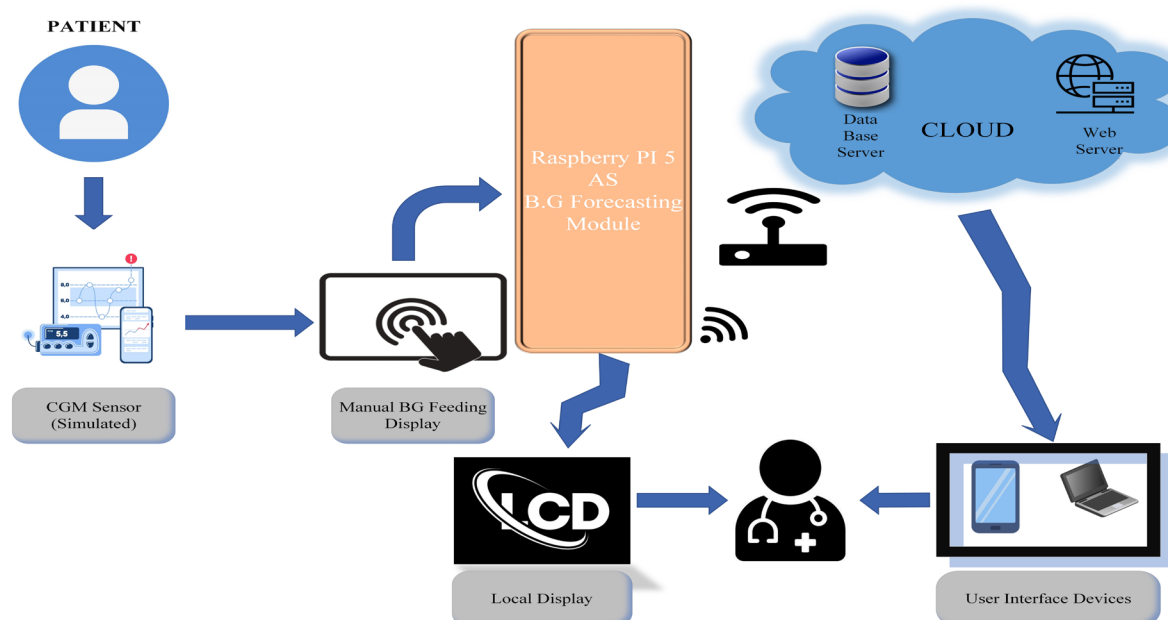
Figure 1: Functional block diagram of SARIMAX-based blood glucose prediction using simulated hourly data.

The block diagram of the operational prototype system is shown in Figure 1. The patient, a Raspberry Pi 5 forecasting module, a manual blood glucose (BG) input interface, a simulated continuous glucose monitoring (CGM) sensor, cloud-based storage, and user notification systems make up the five main parts of the architecture. Although a physical CGM device is not integrated in this prototype, glucose data is simulated at hourly intervals and entered manually via a keypad and display module to emulate real-time CGM behavior. The Raspberry Pi 5 processes the input using a SARIMAX-based forecasting model and visualizes output locally and remotely.

Data Simulation & Pre-Processing

To evaluate forecasting performance, a synthetic dataset of hourly glucose readings was generated spanning a 3-month period (January to March 2024). The dataset emulates realistic CGM behavior, including postprandial spikes and nighttime troughs. Missing entries were handled using linear interpolation, and the dataset was stored in CSV format. Timestamps were converted to datetime objects and aligned to an hourly frequency using the Pandas library. This structured input enabled smooth integration with SARIMAX time-series modeling in Python.

Sarimax Model Configuration

The core of the forecasting system is a Seasonal Auto Regressive Integrated Moving Average with eXogenous variables (SARIMAX) model implemented using the statsmodels package in Python. The

configuration was carefully selected based on autocorrelation and partial autocorrelation plots, yielding:

Non-seasonal order: (p=1, d=1, q=1)
Seasonal order: (P=1, D=1, Q=1, s=24)

These parameters are designed to model daily patterns and seasonality, which are common in blood glucose fluctuations. The Raspberry Pi 5 executes this model to perform hourly forecasts.

Forecasting was carried out for a 7-day horizon (168 steps), starting from March 25 to March 31, 2024. The model was trained on data from January 1 to March 24. To improve clinical interpretability, predicted glucose values were clipped to stay within a physiologically relevant range (70–250 mg/dL). An 80% confidence interval was computed to represent uncertainty bands around the forecasts.

Real-Time Forecasting Execution On Raspberry Pi

The SARIMAX model was deployed on Raspberry Pi 5, which serves as the edge forecasting unit. At runtime, the Pi takes each hourly BG input (simulating CGM readings), appends it to the existing dataset, and executes a rolling forecast using the latest data. The forecasted glucose values are immediately available for local display and cloud synchronization.

A local LCD shows the forecasted value along with a color-coded risk zone (e.g., green for normal, yellow for pre-alert, red for risk). In parallel, the predicted values are transmitted via Wi-Fi using an IP router to a cloud server (such as AWS or Firebase). VISUALIZATION AND PLOTTING

To support analysis and validation of the SARIMAX forecasting model, several plots were generated using matplotlib in Python. The visualization process was divided into three categories

Monthly Historical Plots

Blood glucose values from January 1 to March 31, 2024, were visualized in three separate plots, one per month. These graphs help identify underlying seasonal and daily glucose patterns, and ensure data quality before modeling.

Forecast Plot (Mar 25-31, 2024)

Forecasted glucose levels for the final week of March were visualized with:

- Predicted values (black line)
- 80% confidence intervals (gray band)
- Risk regions shaded (yellow for hypoglycemia, green for normal, red for hyperglycemia)

Actual VS Forecasted Comparison

A comparison plot was also drawn to compare how the actual glucose levels overlap with actual glucose levels of the predicted period of time in which the forecast is made. This is visual evidence on how the model performs particularly in determining glycemic excursions and trends.

Performance Evaluation

Three conventional time-series error measures were used to measure the performance of SARIMAX forecasting model in this study, and they include Mean Absolute Error (MAE), Root Mean Squared

Error (RMSE), and Average Deviation in mg/dL. Such measures are frequently used in the experimental blood glucose forecasting regimes using a CGM because they provide an indication of overall error and because they penalize large error values. The model in this study got the following performance:

- **MAE**: 3.36%
- **RMSE**: 4.30%
- **Average Deviation**: ~4.83 mg/dL

These rates show high rate of forecasting accuracy within the 7-days forecast period using the modeled CGM input. To cite the practice, RMSE values of past SARIMA and SARIMAX-based models in glucose prediction date back to between 5 and 12 mg/dL in numerous studies, which varied depending on the quality of data and the forecasting period [17].

Regarding clinical applicability, the prediction systems are preferably expected to provide foreiteration out of acceptable error rate envisaged by the regulatory framework. FDA guidelines on self-monitoring blood glucose (SMBG) systems indicates that 95 out of 100 records of glucose results must be within + 15 percent and -15 percent of a reference comparator and 99 out of 100 records must be within + 20 percent and - 20 percent in spite of the range of glucose [18]. Although the forecasting model is not a hospitalization diagnostic machine, the error limits finds it well within these limits indicating that it could be very useful in decision-support applications and early warning schemes.

FORECASTING RESULTS AND EVALUATION

The effectiveness of the suggested SARIMAX-based forecasting system of blood glucose was evaluated on the simulated CGM dataset covering three months. A set of visualizations and statistical measures was used to comprehend the adequacy of the model in reflecting the physiological dynamics of glucose. These are a comparison of previous trends in glucose levels, analysis of forecasts in a desired prediction window, and a direct comparison of measured values and forecasted values.The results were interpreted not only in terms of numerical accuracy (using MAE, RMSE, and average deviation) but also by assessing how well the model adheres to clinically relevant glucose ranges. Through these evaluations, the practical applicability and robustness of the SARIMAX model deployed on a Raspberry Pi environment are demonstrated.

Structure and Statistical Overview of Simulated CGM Data Set

Synthetic dataset simulating continuous glucose monitoring (CGM) data was generated to reflect realistic physiological blood glucose (BG) trends over a three-month period from January 1 to March 31, 2024. The dataset contains hourly BG readings, yielding a total of 2,184 data points. These values were carefully designed to emulate typical glycemic patterns in individuals with diabetes, including daily postprandial spikes, overnight glucose dips, and moderate fluctuations based on time-of-day metabolic variation.

Figure 2: Monthly BG Trends from Simulated CGM Data (Jan–Mar 2024)

As shown in Figure 2, the plots across all three months exhibit consistent daily rhythms in glucose levels, with recurring postprandial peaks and nighttime lows. Most glucose readings fall within the target range of 70–180 mg/dL, although occasional excursions into hypo- or hyperglycemic zones were introduced to mimic real-world diabetic variability.

The simulation was developed using Python 3.10 in Visual Studio Code on a Windows 10 system. Key Python libraries included NumPy and Pandas for data generation and structuring, while Matplotlib was used for visualizing the time-series data. The glucose values were synthetically modeled to mirror real CGM behavior, with specific attention paid to meal-associated glucose elevations and circadian drops during sleep hours. The final dataset was saved in CSV format and used both for training the SARIMAX forecasting model and for visual validation.

To verify the integrity and seasonal structure of the dataset, hourly glucose values were plotted month-wise for January, February, and March. Such monthly diagrams provide ultra-fine details of near future/long run glucose dynamics and can be used as a visual starting point on choosing suitable forecasting parameters (e.g. in SARIMAX the seasonal term can set to s=24). There is shading on the glycemic zone on each monthly plot to show clinical cut-offs including yellow (hypoglycemia <70 mg/dL), green (normal range 70 180 mg/dL), and red (hyperglycemia >180 mg/dL). Furthermore, the markers of the color-coded meals are shown at the bottom of the plots that track the simulation of the breakfast, lunch, and dinner events.

Forecast Generation using Sarimax

As the next step to assess the predictive performance of the SARIMAX model, one-week forecasting was conducted based on the synthetic CGM dataset given in Section 4.1. The model has been trained on the data of hourly blood glucose levels during the period of January 1 through March 24, 2024, and the results were used to predict values during the following 7 days (March 25 through March 31, 2024). Through this arrangement, the extent to which the model reflects seasonality and autocorrelation of glucose behavior was estimated especially via daily postprandial cycles and night dips due to glucose that is found in patients with diabetes. The SARIMAX structure to be applied during training was found out through the observations on autocorrelation function (ACF) and partial autocorrelation function (PACF) diagrams, as well as through trial and error.The model parameters selected were:

- **Non-seasonal order:** (p=1, d=1, q=1)
- **Seasonal order:** (P=1, D=1, Q=1, s=24)

This seasonal configuration reflects a 24-hour cycle, consistent with the physiological rhythm of glucose levels throughout the day. The model was implemented in Python using the statsmodels package and executed on a Raspberry Pi 5 as the edge forecasting unit. Upon each hourly input, the device updates the dataset and performs rolling forecasts in real time.

Figure 3: Forecasted Blood Glucose Levels for March 25–31, 2024

As shown in Figure 3, the forecasted glucose levels for the week of March 25–31 follow a smooth, periodic pattern that mirrors the learned trends from the training data. The forecast curve (black line)

remains mostly within the clinically normal range (70–180 mg/dL), with only mild deviations. The 80% confidence interval (gray band) reflects the prediction uncertainty, widening slightly at certain times due to natural fluctuations in the training data. Glycemic risk zones are highlighted using shaded regions: yellow for hypoglycemia (<70 mg/dL), green for normoglycemia (70–180 mg/dL), and red for hyperglycemia (>180 mg/dL).

The forecast visualization demonstrates the SARIMAX model's ability to produce stable and clinically relevant predictions over a medium-range horizon while preserving the temporal structure of the original dataset. This performance highlights its relevance for integration in real-time diabetic monitoring and decision-support systems.

Evaluation of Forecast Accuracy

The competitive expression of the calculations of this figure of the actual compared with the predictive values of the blood glucose (BG) values has been discussed. The main validation of the SARIMAX model is the forecasting period of March 25 31, 2024. This assessment can be quite essential to measure the models statistical performance as well as evaluate the model. a possible clinical use in relative glycemic variation decision support.

Figure 4: Comparison of Actual and Forecasted BG Levels with Error Percentage

In Figure 4 actual glucose data (green line) is overlaid and SARIMAX predicted data. The percentage error is depicted by a red dotted line and the values of the calculation are indicated by a black line; it is done at each hour. Visually, the forecast has a tight bound with the actual data on the 7-day window, as indicated by the strength in the temporal alignment both of amplitude and pattern. The SARIMAX model takes into account the daily glucose levels perfectly well. such things as postprandial peaks, evening plunges, and intermeal plateau periods-in imitation of an efficient learning of seasonal patterns of the time series.

The red dotted line gives a real time indication of the discrepancy between the actual and the predicted indicator, which in turn presents how and when the model is a little over- or underestimating

the presence of glucose in the body. The majority of the percentages of error are not that high but still experience a spike around the peak changes or other points where the data begins to bend approximately - peak transitions or inflection points - which is characteristic of physiological signals and where there is a dramatic change in glucose levels due to simulated meals or the effects of insulin. Such deviations could be called non-clinical and the fact is that most of the predictions stay within the acceptable range of glucose 70180 mg/dL.

Notably, the forecast has no cumulative drift or destabilization in the course of time. verifying the temporal stability and generalization of SARIMAX model. Such uniformity is riotous in real-time applications, where models have to be changed continually with no loss in quality accuracy. This kind of behavior has been similarly reported elsewhere under comparative studies where traditional ARIMA-based models are used. had good short term performance but consistent error profiles [19].

Whereas superior models such as LSTM and GRU have demonstrated good predictive powers, they are typically laborious in their needs in terms of datasets and interpretability, and they are also costly in terms of computations. requirements, particularly when they are used on edge platforms [20]. On the contrary, SARIMAX proposes a desirable trade-off among accuracy, interpretability and the capability of processing. More recent interpretable is also hot on the topic of emergent visibility of algorithms in clinical time-series makes the forecasting tasks [21].

CONCLUSION

This study presents a SARIMAX-based blood glucose forecasting system that demonstrates high accuracy, transparency, and practical deployment capability using a Raspberry Pi 5 edge-computing setup. The approach leverages synthetically generated CGM data designed to mimic realistic glycemic patterns observed in diabetic patients. Through structured time-series modeling, the SARIMAX model effectively captures daily seasonal fluctuations and postprandial dynamics, producing clinically relevant glucose forecasts over a 7-day horizon.

Key evaluation metrics, namely, a Mean Absolute Error of 3.36%, RMSE of 4.30%, and an average deviation of approximately 4.83 mg/dL, confirm the model's predictive strength. Visualization of historical trends, forecasted values, and actual versus predicted comparisons further validates the temporal consistency and reliability of the forecasting system. The integration of color-coded glycemic risk zones and real-time display on Raspberry Pi enhances the system's utility for patient monitoring and early intervention.

Unlike black-box machine learning models, SARIMAX offers interpretable forecasting with minimal computational overhead, making it suitable for real-time applications and edge environments. The model adheres well to regulatory accuracy thresholds and shows potential for use in decision-support systems for diabetes care. Future work may focus on incorporating real-world CGM data, expanding the forecasting horizon, and integrating external variables such as physical activity or meal content to further improve model performance and personalization.

REFERENCES

International Diabetes Federation, "IDF Diabetes Atlas," 10th ed., 2021. [Online]. Available: https://diabetesatlas.org (Accessed: Jul. 10, 2025)

American Diabetes Association, "Standards of Medical Care in Diabetes—2023," *Diabetes Care*, vol. 46, suppl. 1, pp. S1–S291, Jan. 2023. [Online]. Available: https://doi.org/10.2337/dc23-S001(Accessed: Jul. 10, 2025)

J.-L. Díez, E. Montaser, and J. Bondia, "Glucose prediction under variable-length time-stamped daily events: A seasonal stochastic local modeling framework," *Sensors*, vol. 21, no. 9, p. 3188, May 2021. [Online]. Available: https://doi.org/10.3390/s21093188(Accessed: Jul. 10, 2025)

C. Zecchin et al., "A review on the use of prediction models in diabetes management," *IEEE Rev. Biomed. Eng.*, vol. 14, pp. 287–302, 2021.

M. Anjum, R. Saher, and M. N. Saeed, "AI-enhanced time series analysis of continuous glucose monitoring data for personalized dietary intervention," *PeerJ Comput. Sci.*, vol. 10, e1971, Apr. 2024. [Online]. Available: https://doi.org/10.7717/peerj-cs.1971(Accessed: Jul. 10, 2025)

S. P. Shashikumar et al., "Artificial intelligence in diabetes: A systematic review," *NPJ Digit. Med.*, vol. 5, no. 1, pp. 1–11, 2022.

A. Kistkins, T. Zhang, and S. Wang, "Comparative analysis of ARIMA, LSTM, and logistic regression for blood glucose prediction," *BMC Med. Inform. Decis. Mak.*, vol. 23, no. 1, 2023. [Online]. Available: https://doi.org/10.1186/s12911-023-02111-7(Accessed: Jul. 10, 2025)

A. Contreras, J. Vélez, and J. Bondia, "An interpretable SARIMA-based model for short-term blood glucose prediction using real-world CGM data," *Frontiers in Artificial Intelligence*, vol. 6, p. 1182127, 2023.

C. Barakat, A. Fakih, and H. Hajj, "Real-time glucose forecasting models for wearable systems: A review," *IEEE Access*, vol. 9, pp. 154451–154471, 2021.

A. Facchinetti et al., "Real-time glucose forecasting in type 1 diabetes: Use cases and challenges," *J. Diabetes Sci. Technol.*, vol. 15, no. 3, pp. 588–599, May 2021.

A. Daskalaki et al., "Glucose prediction using event-based time-series modeling in type 1 diabetes," *IEEE Trans. Biomed. Eng.*, vol. 68, no. 2, pp. 638–648, Feb. 2021.

D. Lu, X. Liu, and M. C. Cheung, "Simulation-based CGM datasets for training and evaluation of forecasting models," *Healthcare Analytics*, vol. 2, 2022, Art. no. 100073.

C. Sevil, G. Quintana-Ortí, and J. Bondia, "A systematic review of prediction models for glycemia forecasting in diabetes management," *Sensors*, vol. 22, no. 6, p. 2205, Mar. 2022. [Online]. Available: https://doi.org/10.3390/s22062205(Accessed: Jul. 10, 2025)

L. D. Moser et al., "Enhancing CGM forecasting accuracy with clustering-enhanced SARIMA," *J. Diabetes Sci. Technol.*, vol. 16, no. 3, pp. 505–514, 2022. [Online]. Available: https://doi.org/10.1177/19322968211039547(Accessed: Jul. 10, 2025)

J. Bondia et al., "Trends and innovations in glucose prediction algorithms for diabetes management," *Sensors*, vol. 22, no. 10, p. 3952, 2022. [Online]. Available: https://doi.org/10.3390/s22103952(Accessed: Jul. 10, 2025)

L. B. Katz, L. Stewart, D. King, and H. Cameron, "Meeting the new FDA standard for accuracy of self-monitoring blood glucose test systems intended for home use by lay users," *J. Diabetes Sci. Technol.*, vol. 14, no. 3, pp. 580–586, May 2020. [Online]. Available: https://doi.org/10.1177/1932296820906184(Accessed: Jul. 10, 2025)

M. Gadaleta, F. Facchinetti, A. Sparacino, and C. Cobelli, "Glucose forecasting and risk evaluation for intelligent diabetes management," *IEEE Transactions on Biomedical Engineering*, vol. 69, no. 1, pp. 122–130, Jan. 2022.

J. Bondia et al., "Hybrid models for glycemia forecasting: Evaluating SARIMA and AI integration," *Sensors*, vol. 22, no. 6, p. 3952, 2022.

A. Mirshekarian, H. Nazeran, and S. Tulu, "Forecasting blood glucose levels using ARIMA and LSTM models: A comparative analysis," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 1, pp. 81–90, Jan. 2023.

E. Georga, V. Protopappas, and D. Fotiadis, "Glucose prediction in type 1 and type 2 diabetes using machine learning and stochastic modeling: State-of-the-art and future directions," *Computers in Biology and Medicine*, vol. 152, 2023, Art. no. 106388. [Online]. Available: https://doi.org/10.1016/j.compbiomed.2022.106388(Accessed: Jul. 10, 2025)

F. Jiang, C. Wang, and Y. Shen, "Lightweight interpretable forecasting models for blood glucose prediction using time series clustering and ARIMA variants," *Artificial Intelligence in Medicine*, vol. 130, 2022, Art. no. 102319.