JOURNAL OF
EMERGING INVESTIGATORS

Article

Using artificial intelligence to forecast continuous glucose monitor (CGM) readings for type one diabetes

Ishita Jalla[1], Sindhu Ghanta[2]

[1] Eastlake High School, Sammamish, Washington

[2] AIClub, Mountain View, California

SUMMARY

Type One diabetes (T1D) is an incurable condition in which a person produces little to no insulin, leading to high glucose levels in the bloodstream. Advancements in technology have led to the development of Continuous Glucose Monitors (CGM), involving inserting a sensor under the skin that continuously measures glucose levels in the bloodstream. CGMs provide real-time data on glucose levels and patterns of fluctuation and have evolved as an essential part of diabetes management. Current CGMs cannot forecast glucose levels. In this study, we hypothesized that machine learning algorithms can be utilized to forecast glucose values using data provided by the CGM device. This could help T1D patients optimize medication timing, dietary adjustments, and physical activity to improve metabolic control and prevent T1D complications. Our study used publicly available data from a study that compared CGMs with and without routine blood glucose monitoring in adults with T1D. We used different AI models to provide 30-minute forecasts of blood glucose levels to show how model structure and parameters affect Root Mean Squared Error (RMSE), and whether the error can be reduced to 5% as per medical standards. Our research used RandomForest as a baseline algorithm and LSTM, a recurrent neural network. Our results demonstrated that the LSTM model achieved an error of under 5% in predicting future blood glucose levels. Hence, researchers can use LSTM models to enhance CGMs to aid T1D patients. The results of our study supported our hypothesis that AI can be used to predict glucose levels.

INTRODUCTION

In 2017, there were 9 million recorded people with Type One Diabetes (T1D) (1). T1D is a condition in which the pancreas produces little to no insulin. Without insulin, blood sugar can't get into cells and builds up in the bloodstream (1). Over time high glucose levels in the bloodstream will cause damage to the body. The millions of people who combat diabetes need to control aspects of their lifestyle, such as food and sleep, to keep their blood sugar under control (1). Unfortunately, there is no cure. As advancements in technology progress, new methods for addressing diabetes have emerged. One such innovation is the Continuous Glucose Monitor (CGM). CGMs involve inserting a sensor under the skin, typically in the arm or abdomen, which continuously measures glucose levels in the bloodstream (2). These devices provide real-time data on glucose levels and fluctuation patterns throughout the day, making them a crucial tool for individuals living with diabetes to manage their condition (2). In the past decade, continuous glucose monitoring has become an essential part of diabetes management for many people with T1D (3).

Before CGMs, blood glucose levels in diabetes patients were monitored through self-monitoring of blood glucose (SMBG). However, SMBG has problems such as user error in test accuracy, the need for multiple finger-blood stick samples daily, and the limited amount of data SMBG provides (4). The utilization of implantable glucose sensors has been a recognized concept for the past four decades (5). However, a recent study has demonstrated that CGM improves the quality of life for individuals with diabetes by providing more optimized glucose control and facilitating better disease management and the use of CGM is as effective as testing blood glucose levels in a lab (3, 5).

A CGM helps people with diabetes better understand and manage their conditions. By analyzing the trends and patterns in their blood sugar levels over time, they can identify factors that might be affecting their glucose levels and take the correct steps to address them (2). The use of CGM is associated with improvement in metabolic control in T1D and has been demonstrated to improve metabolic control in individuals with T1D (2). However, to further enhance blood sugar management in this population, the capability of CGM devices to predict future blood glucose levels is crucial. The ability to forecast glucose levels offers numerous benefits for patients, including preventing erratic fluctuations and making more informed decisions regarding treatment (5). Machine learning algorithms can be utilized to forecast glucose values using historical data provided by the CGM device (5). These predictions can aid in optimizing medication timing, dietary adjustments, and physical activity to ultimately improve metabolic control and prevent complications associated with T1D.

As a next step, researchers are looking into building intelligent models that can predict future glucose levels based on historical information. Predicting future blood glucose levels is vital because it can help individuals with diabetes manage their condition more effectively. Through artificial intelligence, scientific researchers have been trying to make an accurate prediction method for blood glucose levels (6). Applying machine learning and artificial intelligence (AI) techniques to forecast glucose levels in individuals with diabetes is a relatively under-explored field. While a prior study has employed various inputs in addition to CGM data, the primary focus of this research is the utilization of CGM data as the sole input (6). This is due to the accessibility and convenience of CGM for diabetes patients. Some studies

https://doi.org/10.59720/23-291

employing CGM data as inputs utilized different algorithms than those employed in this study.

We used two machine-learning models in our study: Long Short-Term Memory Networks (LTSM), which is a recurrent neuronal network, and RandomForest. A deep learning model comprises various layers of nodes, including an input layer, one or multiple hidden layers, and an output layer. Recurrent neural networks (RNNs) employ previous outputs as inputs and possess hidden states, thereby enabling the usage of causal information and hidden patterns in time-series data. However, RNNs suffer from not being able to model long sequences. Long Short-Term Memory Networks (LSTMs), were first introduced by Hochreiter and Schmidhuber in 1997 and have been widely used in numerous applications such as natural language processing, speech recognition, and time series forecasting (7). LSTM is well-suited for sequential data such as time series. LSTMs can be calibrated by adjusting the learning rate and the number of epochs (8). The number of epochs determines how many times the weights of the neural network are changed, and the learning rate determines the pace at which the algorithm learns (8).

RandomForest is a commonly used machine learning algorithm used for classification and regression, trademarked by Leo Breiman and Adele Cutler in 2001, which combines the output of multiple decision trees to reach a single result (9). The decision tree starts with one basic question. Based on that answer to the base question, other questions get formulated, and the answers to those questions make up decision nodes in the trees (9). The main components of the random forest algorithm are the nodes in the trees, the number of trees, and the number of features sampled (9).

In this study, we hypothesized that RandomForest and LSTM techniques will be effective in forecasting the blood glucose levels of a patient 30 minutes ahead, with a margin of root mean square error not exceeding five units. We experimented with several machine learning and deep learning algorithms such as RandomForest and Long-Short Term Memory (LSTM) and identified a model able to predict glucose levels within an acceptable range of error for a medical device. Sixty epochs with a learning rate of 0.0001 LSTM units performed the best on the test data with an RMSE value of 4.8.

RESULTS

RandomForest Model

We used two AI models. The first was RandomForest, which we ran on Amazon Web Services. We wanted to start with a simple regression model and decided on RandomForest because it is effective with large quantities of data (9). Our goal was to reach an RMSE value below 5. The hyperparameters that were adjusted in the RandomForest Model to reach this goal were the number of trees and the depth. The number of trees varied from 10, 20, 30, 40, 50, 60, 80, 100, and 150. The depth varied from 1, 2, 3, 4, 5, or 6. The results for the lowest RMSE value we got from RandomForest was 28.4844 when the number of trees was 150 and the depth was 6 (number of trees = 150, depth = 6) (**Table 1, Figure 1**).

LSTM Model

After establishing a baseline with RandomForest, we ran LSTM models on Google Colab with the same goal of reaching an RMSE below 5. The advantage of LSTM models over RandomForest is in their ability to handle time series data because they can remember earlier time steps which yields more accurate predictions (8). Through tuning hyperparameters, we got many results with an RMSE value below 5. The hyperparameters that were adjusted were the epochs and the learning rate. The epochs varied from 10, 20, 30, 40, 50, 60, 80, 100, and 150 and the learning rate varied from 0.00001, 0.0005, 0.0001, 0.005, 0.001, and 0.05 (**Table 2, Figure 2**). The lowest RMSE value was 4.8118 when the number of epochs was 60 and the learning rate was 0.001. The most dramatically different model was with a learning rate of 0.00001: this model yielded RMSE values that were way higher than the other LSTM models produced (**Table 2, Figure 2**). All models except for those with a learning rate of 0.00001 and 0.005 were similar in the range of RMSE values they produced (**Table 2, Figure 3**).

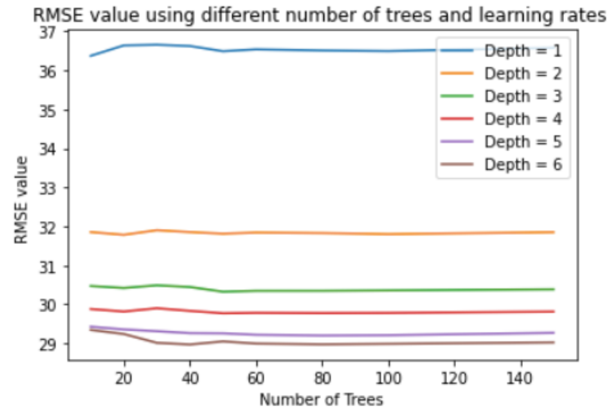

Figure 1: RMSE values of the RandomForest experiments. Change of RMSE values based on the depth (how far the tree goes) and number of trees using the RandomForest model. The lower the RMSE value the better the results. As the trees' depth increased, the RMSE value decreased.

DISCUSSION

The goal of our research was to forecast blood glucose

		Depth					
		1	2	3	4	5	6
Number of Trees	10	33.6375	31.0835	30.3489	29.8625	29.5173	28.861
	20	33.6145	31.0937	30.344	29.8973	29.2987	28.7138
	30	33.3669	31.1008	30.3927	29.8315	29.1436	28.5772
	40	33.4135	31.0534	30.3585	29.794	29.0821	28.6377
	50	33.4167	31.0539	30.3251	29.734	29.0121	28.576
	60	33.3924	31.0675	30.3292	29.7948	29.1158	28.6161
	80	33.4369	31.0468	30.263	29.7931	29.1632	28.5284
	100	33.3999	31.059	30.2515	29.7524	29.1295	28.5093
	150	33.4176	31.0738	30.2494	29.7508	29.1123	28.4844

Table 1: Results of the RandomForest Experiments. RMSE yielded from the RandomForest Model based on the number of trees and depth hyperparameters.

https://doi.org/10.59720/23-291

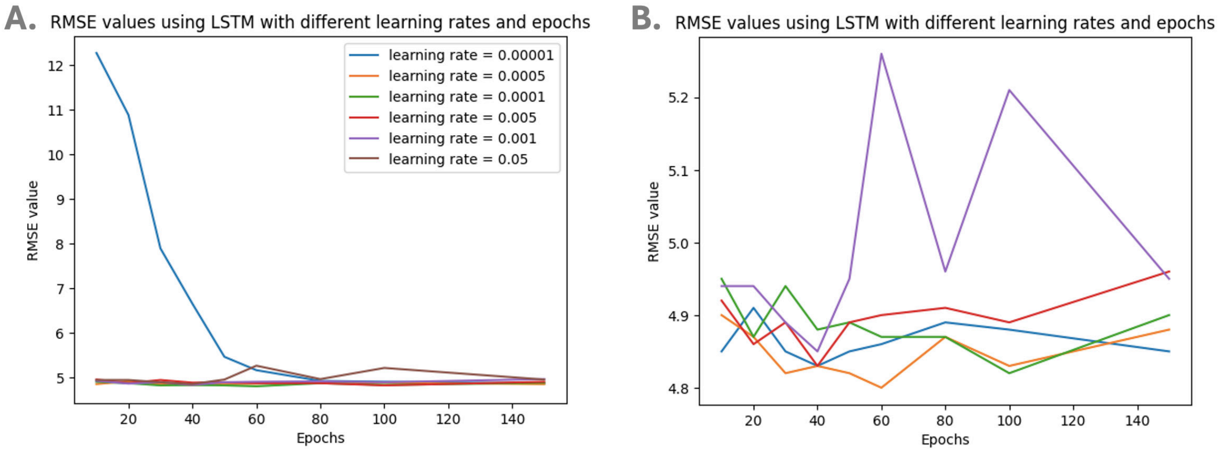

Figure 2: RMSE values using the LSTM Model with different learning rates and epochs. RMSE values change when the learning rate (how fast the model learns) and the epochs hyper-parameters are changed. RMSE values A) with and B) without the learning rate of 0.00001. B) Jumps are clearer without 0.00001 showing distinct learning rates.

readings thirty minutes into the future based on historical data by exploring different machine learning and deep learning algorithms. A study conducted in July 2020 utilized similar algorithms (LSTM), to those used in this work, but only included data from 10 patients while our study includes data from 225 patients incorporating a larger patient population leads to more diverse and accurate data, which in turn improves the performance of deep learning or machine learning algorithms (6, 7). We transformed the time-series nature of the CGM data into a regression problem and implemented statistical machine-learning models, such as RandomForest. We used RMSE to evaluate the performance of the algorithm. The results of the Random Forest experiments were not promising, and all of the RMSE values were unacceptable. The hyperparameters that we tuned were the number of trees and the depth. We used this algorithm as a baseline and then moved on to the second algorithm, LSTM. LSTM was employed for forecasting CGM values. In the case of LSTM, the data did not need to be converted into a regression problem, as LSTM architecture inherently consists of a forecasting component.

The glucose level forecasting problem was converted into a regression problem to enable the use of statistical machine-learning algorithms. The Random Forest algorithm, which is a type of ensemble learning method, was used to predict glucose levels. The RMSE, a measure of the difference between the predicted values and the actual values, was used to evaluate the performance of the algorithm. A study published by the Journal of Diabetes Science and Technology established that predicting glucose levels up to 5 of error is medically acceptable (7). With this basis, the goal of our research was to explore machine learning algorithms that can forecast blood glucose levels based on historical data. The results showed that the maximum RMSE values obtained were 33.63, while the minimum RMSE values obtained were 28.4844. These values fall outside the acceptable range of error, an RMSE of 5 (7), for a medical application such as this. For RandomForest the hyperparameters we tuned were the number of trees and depth.

Further analysis revealed that the depth of the algorithm had a high impact on improving its performance of the algorithm. Increasing the depth of the algorithm resulted in a

		Learning Rate					
		0.00001	0.0005	0.0001	0.005	0.001	0.05
Epochs	10	12.27	4.85	4.90	4.95	4.92	4.94
	20	10.88	4.91	4.87	4.87	4.86	4.94
	30	7.89	4.85	4.82	4.94	4.89	4.89
	40	6.65	4.83	4.83	4.88	4.83	4.85
	50	5.46	4.85	4.82	4.89	4.89	4.95
	60	5.16	4.86	4.80	4.87	4.90	5.26
	80	4.92	4.89	4.87	4.87	4.91	4.96
	100	4.90	4.88	4.83	4.82	4.89	5.21
	150	4.86	4.85	4.88	4.90	4.96	4.95

Table 2: RMSE values produced from LSTM models. RMSE values yielded through the LSTM model with various epochs and learning rates.

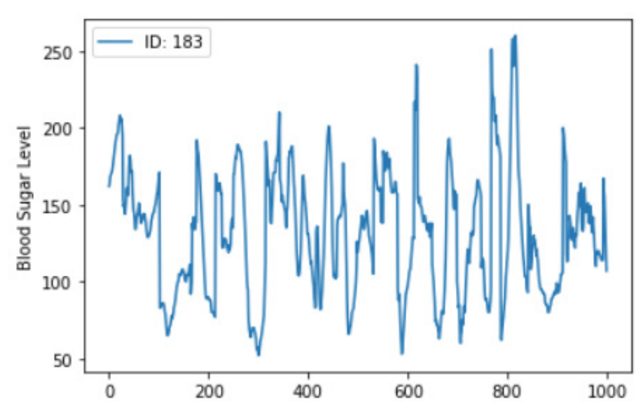

Figure 3: Fluctuation of blood glucose level. Fluctuation of patient #183's blood glucose level over six months.

https://doi.org/10.59720/23-291

Hour	Minute	F1	F2	F3	F4	Prediction
5	35	162mg/dL	164mg/dL	168mg/dL	169mg/dL	180mg/dL
5	30	164mg/dL	168mg/dL	169mg/dL	170mg/dL	183mg/dL
5	25	168mg/dL	169mg/dL	170mg/dL	171mg/dL	186mg/dL
5	20	169mg/dL	170mg/dL	171mg/dL	172mg/dL	190mg/dL
5	15	170mg/dL	171mg/dL	172mg/dL	174mg/dL	191mg/dL
5	10	171mg/dL	172mg/dL	174mg/dL	177mg/dL	193mg/dL
5	5	172mg/dL	174mg/dL	177mg/dL	180mg/dL	195mg/dL
5	0	174mg/dL	177mg/dL	180mg/dL	183mg/dL	196mg/dL
4	55	177mg/dL	180mg/dL	183mg/dL	186mg/dL	196mg/dL
4	50	180mg/dL	183mg/dL	186mg/dL	190mg/dL	198mg/dL
4	45	183mg/dL	186mg/dL	190mg/dL	191mg/dL	202mg/dL

Table 3: Processed data with each separate IDs. The dataset that went through processing has properties of hour, minute, F1 (glucose at the time), F2 (glucose value after a minute), F3 (glucose value after two minutes), F4 (glucose value after three minutes), and the predicted glucose value.

decrease in the RMSE values. On the other hand, increasing the number of trees in the algorithm did not have a notable impact on improving the performance of the algorithm. This meant that the RMSE values yielded were still way too high to be medically acceptable.

The LSTM model, on the other hand, had a much lower RMSE compared to the RMSE of RandomForest. The LSTM model we used comprises a single-layer LSTM with 50 units. Each unit of the LSTM layer is responsible for memorizing and updating information from the previous time steps(10). The model has 10,451 trainable parameters, which are the weights and biases that are adjusted during training to optimize the model's performance. The output of the model is a single number that represents the predicted glucose level in the next 30 minutes. The input to the model is the glucose level in the past one and a half hours. By providing the glucose level at previous time steps, the model can take into account the temporal dynamics of the data and make more accurate predictions. The results showed that the maximum RMSE values obtained were 12.53, while the minimum RMSE values obtained were 4.8118. These values fall within an acceptable margin of error for a medical application such as this. The hyper-parameter learning rate had a high impact on the model performance. Increasing the number of epochs on the other hand, did not help improve the performance of the model, which shows that the model converged in a small number of epochs. Having a learning rate that was too small caused the model to converge very slowly. Having a large learning rate, on the other hand, made the model jump around a lot in accuracy. The lowest RMSE value, 4.80, was obtained by LSTM for a learning rate of 0.0001 and an epoch value of 60.

The LSTM model had a much lower RMSE compared to the RMSE of the Random Forest model, and the values fall within an acceptable range of error for a medical application. The lowest RMSE value, 4.80, was obtained by LSTM for a learning rate of 0.0001 and an epoch value of 30. Based on our study, LSTM models would be suitable for predicting glucose readings in CGM devices to aid T1D patients in preventing the harmful effect of erratic trends in glucose levels.

MATERIALS AND METHODS

The software libraries used were Pandas Version 2.2.1., NumPy Version 1.26.4, Sklearn Version 1.4.1, and TensorFlow 2.15.0. We used TensorFlow for running deep-learning RNNs. Google Colab was used to code everything and run experiments. Machine learning algorithms (RandomForest) were run using Amazon Web Services.

Data Availability and Acquisition

The training data was obtained from adults aged 25-60 with T1D from a study by the American Diabetes Association that determined whether using CGMs was as effective as using blood glucose monitoring (2). The total number of participants was 225. The CGM used was the Dexcom G4 which monitored the patients for six months. Data was collected every five minutes, and over ten thousand samples were collected. Further data processing was done to convert the dataset into a regression problem. Converting the dataset allowed the use of a statistical machine learning algorithm like a random forest to make predictions. To convert the dataset, we took each patient's data and converted it from a dataset that had ID, time, glucose level, and hour and minute columns to a dataset that had columns of the hour and minute, four features, and the prediction (Table 3). The original dataset contained all the IDs in a sequence but for converting the dataset we had to process each ID separately. The interval between the features was one, and the interval between the first feature and the prediction was nine.

Machine Learning Analysis

RandomForest was run with the depth was fixed at 1, 2, 3, 4, 5, or 6. The number of trees was fixed at either 10, 20, 30, 40, 50, 60, 80, 100, or 150. The second model, LSTM, was also trained with two pre-determined hyperparameters. These were epochs and learning rates. The epochs were fixed at either 10, 20, 30, 40, 50, 60, 80, 100, or 150 and the learning rate varied from 0.00001, 0.0005, 0.0001, 0.005, 0.001, and 0.05. These models yielded an RMSE value which was calculated by taking the square root of the difference between the predicted value and the observed value and then squaring the resulting value.

Received: October 12, 2023
Accepted: January 24, 2024
Published: August 7, 2024

REFERENCES

1. "Diabetes." *World Health Organization.* www.who.int/news-room/fact-sheets/detail/diabetes. Accessed October 25, 2023.
2. Patton, Susana R, et al. "Continuous Glucose Monitoring Versus Self-monitoring of Blood Glucose in Children with Type 1 Diabetes- Are there Pros and Cons for Both?" *US Endocrinology*, vol. 8, no. 1, 2012, pp. 27-29. PMCID: PMC3848052
3. Floyd, Baraka, et al. "Comparative Analysis of the Efficacy of Continuous Glucose Monitoring and Self-Monitoring of Blood Glucose in Type 1 Diabetes Mellitus." *Journal of diabetes science and technology*, vol. 6, no. 5, 09/12/2012, pp. 1094-102, https://doi.org/10.1177/193229681200600513.
4. Didyuk, Olesya, et al. "Continuous Glucose Monitoring Devices: Past, Present, and Future Focus on the

https://doi.org/10.59720/23-291

History and Evolution of Technological Innovation." *Journal of diabetes science and technology*, vol. 15, no. 3, 01/13/2020, pp. 676–683, https://doi.org/10.1177/1932296819899394.
5. Ali, Jaouher, et al. "Continuous Blood Glucose Level Prediction of Type 1 Diabetes Based on Artificial Neural Network." *Biocybernetics and Biomedical Engineering*, vol. 38, no. 4, 02/26/2021, pp. 828-840, https://doi.org/10.1016/j.bbe.2018.06.005.
6. Idrissi, Touria, et al. "Deep Learning for Blood Glucose Prediction: CNN vs LSTM". *International Conference on Computational Science and Its Applications*, vol. 12250, 08/19/2020, pp. 379-393. https://doi.org/10.1007/978-3-030-58802-1_28.
7. Krouwer, Jan, et al. "A Review of Standards and Statistics Used to Describe Blood Glucose Monitor Performance." *Journal of diabetes science and technology*, vol. 4, no. 1, 01/01/2010, pp. 75–83. https://doi.org/10.1177/193229681000400110.
8. Saxena. "LSTM: Introduction to LSTM: Long Short Term Memory." *Analytics Vidhya*, www.analyticsvidhya.com/blog/2021/03/introduction-to-long-short-term-memory-lstm/. Accessed October 25, 2023.
9. "What is Random Forest?" IBM. www.ibm.com/topics/random-forest#:~:text=IBM,to%20reach%20a%20single%20result. Accessed October 25, 2023. Hochreiter, Sepp, et al. "Long Short-Term Memory." *Neutral computation*, vol. 9, no. 8, 11/15/1997, pp. 1735-1780. https://doi.org/10.1162/neco.1997.9.8.1735.