



NATIONAL TECHNICAL UNIVERSITY OF ATHENS
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING
SCHOOL OF MECHANICAL ENGINEERING

INTERDISCIPLINARY POSTGRADUATE PROGRAMME
“Translational Engineering in Health and Medicine”

Development of attention-based LSTM models for the prediction of nocturnal hypoglycemia in patients with Type 1 Diabetes

Postgraduate Diploma Thesis

Konstantinos Tziavaras

Supervisor: Konstantina Nikita, Professor NTUA

Athens, October 2024



NATIONAL TECHNICAL UNIVERSITY OF ATHENS
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING
SCHOOL OF MECHANICAL ENGINEERING

INTERDISCIPLINARY POSTGRADUATE PROGRAMME
“Translational Engineering in Health and Medicine”

Development of attention-based LSTM models for the prediction of nocturnal hypoglycemia in patients with Type 1 Diabetes

Postgraduate Diploma Thesis

Konstantinos Tziavaras

Supervisor: Konstantina Nikita, Professor NTUA

The postgraduate diploma thesis has been approved by the examination committee on
14/10/24

1st member

.....

Konstantina Nikita,
Professor NTUA

2nd member

.....

Giorgos Stamou,
Professor NTUA

3rd member

.....

Athanasios Voulodimos,
Assistant Professor

Athens, October 2024

.....
Konstantinos Tziavaras

Graduate of the Interdisciplinary Postgraduate Programme,
“Translational Engineering in Health and Medicine”,
Master of Science,
School of Electrical and Computer Engineering,
National Technical University of Athens

Copyright © - *Konstantinos Tziavaras, 2024*

All rights reserved.

You may not copy, reproduce, distribute, publish, display, modify, create derivative works, transmit, or in any way exploit this thesis or part of it for commercial purposes. You may reproduce, store or distribute this thesis for non-profit educational or research purposes, provided that the source is cited, and the present copyright notice is retained, Inquires for commercial use should be addressed to the original author.

The ideas and conclusions presented in this paper are the author’s and do not necessarily reflect the official views of the National Technical University of Athens.

Abstract

Diabetes Mellitus (DM) is a chronic condition with a rising global prevalence and severe complications. The International Diabetes Federation projects that the number of individuals with diabetes will reach 643 million by 2030. To enhance glycemic control and mitigate the risk of serious physical and emotional complications related to hypoglycemia, this thesis presents the design, development, and evaluation of an interpretable model for predicting the risk of nocturnal hypoglycemic episodes in individuals with Type 1 Diabetes (T1DM). The proposed model employs a hybrid approach, integrating compartmental models with machine learning techniques. The OhioT1DM dataset, which includes real data from the eight-week monitoring period of 12 patients, was utilized for both development and evaluation purposes. Input data for the model consisted of glucose measurements, insulin doses, and meal information from the previous 24 hours. Mathematical models for simulating (i) the physiological mechanisms of insulin absorption from the subcutaneous tissue into the bloodstream, (ii) the activation of the insulin signaling pathway, and (iii) the absorption of glucose from the intestine were combined with Long Short-Term Memory Neural Networks (LSTMs). A custom attention layer was integrated to enhance the model's performance and provide insights into the model's reasoning behind its predictions. The model was assessed in terms of its ability to correctly predict nocturnal hypoglycemic events within a twelve-hour prediction window. Moreover, the Monte Carlo Dropout method was applied to quantify the uncertainty of the model's predictions. The model was also evaluated on an external dataset from the ten-day monitoring period of 12 T1DM patients, which was granted from the Diabetes Center, First Department of Pediatrics, P. & A. Kyriakou Children's Hospital, Athens, within the framework of the SMARTDIAB project.

Keywords: Type 1 Diabetes, Machine Learning, Deep Learning, Compartment Models, Hybrid Model, Compartment Models, Risk Prediction Model, Explainable Artificial Intelligence, XAI, Attention, Monte Carlo Dropout

Acknowledgements

I would like to express my heartfelt gratitude to my professor, Konstantina Nikita, for her guidance and for giving me the unique opportunity to work on a project of great research significance and immediate social impact. I am also deeply thankful to postdoctoral researcher Maria Athanasiou, with whom I had the pleasure of collaborating. Her valuable time, insightful feedback, and motivation were instrumental in helping me complete this work and broaden my knowledge.

I am equally grateful to my friends and my girlfriend, Vasia, for their unwavering support, love, understanding, and belief in me.

Finally, I owe a profound thank you to my family. My deepest appreciation goes to my grandmother, Olympia, from whom I inherited the "spark" of curiosity and a passion for research. My younger brother, Antonis, is often the voice of reason when I need it, helping to keep my thoughts grounded. And, of course, my mother, Eleni, without whom I wouldn't have achieved anything in my life.

Contents

Abstract.....	5
Acknowledgements.....	6
List of Figures	11
List of Tables	13
1. Introduction	1
1.1. Research Problem	1
1.2 Objective of the Thesis.....	1
1.3 Structure of the Thesis.....	2
2. Diabetes Mellitus	3
2.1 Glucose Metabolism and Regulation	3
2.1.1 Glucose Metabolism: Sources and Pathways	3
2.1.2 Regulation of Blood Glucose	3
2.1.2.1 Insulin.....	4
2.1.2.2 Glucagon	5
2.1.2.3 Insulin–Glucagon Balance in Blood Glucose Regulation	5
2.2 Diabetes Overview and Types.....	6
2.2.1 Type 1 Diabetes.....	6
2.2.2 Type 2 Diabetes.....	7
2.2.3 Gestational Diabetes Mellitus.....	7
2.3 Diabetes Symptoms and Complications	7
2.3.1 Symptoms	7
2.3.2 Acute Complications	8
2.3.2.1 Diabetic Ketoacidosis (DKA)	8
2.3.2.2 Hyperglycemia Hyperosmolar State (HHS)	9
2.3.2.3 Hypoglycemia.....	9
2.3.2.4 Nocturnal Hypoglycemia.....	9
2.3.3 Long-Term Complications	10
2.3.3.1 Diabetic Retinopathy	11
2.3.3.2 Diabetic Nephropathy.....	12
2.3.3.3 Diabetic Cardiovascular Disease	12
2.3.3.4 Diabetic Neuropathy.....	13
2.4 Diabetes Diagnosis	14

2.4.1 Hemoglobin A1C (HbA1c) Test.....	14
2.4.2 Fasting Plasma Glucose (FPG) Test	14
2.4.3 Glucose Tolerance Test (GTT)	14
2.5 Diabetes Management.....	15
2.5.1 Blood Glucose Monitoring	15
2.5.2 Insulin and Antidiabetic Drugs	16
2.5.3 Nutrition and Exercise.....	17
3. Artificial Neural Networks.....	18
3.1 Emulating Brain Function for Machine Learning	18
3.2 Perceptron	18
3.2.1 Origin and Function.....	18
3.2.2 Fundamentals of the Perceptron	19
3.2.3 Activation Functions	20
3.2.3.1 Threshold Function	20
3.2.3.2 Linear Function	20
3.2.3.3 Sigmoid Function.....	21
3.2.3.4 Hyperbolic Tangent (tanh) Function	21
3.2.3.5 Rectified Linear Unit (ReLU) function	21
3.2.3.6 Softmax Function	21
3.3 Multilayer Perceptrons	22
3.4 Recurrent Neural Networks	24
3.4.1 Long – Short Term Memory Networks.....	25
3.5 Machine Learning Types	28
3.5.1 Supervised learning.....	28
3.5.2 Unsupervised Learning.....	30
3.5.3 Reinforcement Learning.....	31
3.6 Development of ANNs for Supervised Learning tasks	32
3.6.1 Data Preprocessing	32
3.6.1.1 Missing Data Handling	32
3.6.1.2 Data Normalization	33
3.6.2 Training of the Model	34
3.6.3 Evaluation of the Model.....	35
4. Related Work	38

4.1 Physiological Prediction Models	38
4.2 Machine Learning Models for Predicting Blood Glucose Levels	39
4.3 Machine Learning Models for Predicting Nocturnal Hypoglycemia	40
5. Development of attention-based LSTM models for the prediction of nocturnal hypoglycemia events in T1DM.....	43
5.1 Dataset and Data Preparation	43
5.1.1 OhioT1DM Dataset.....	43
5.1.2 Data Preparation.....	44
5.2 Architecture of Hybrid Predictive Model.....	46
5.2.1 Intestinal Glucose Absorption Compartment Model.....	47
5.2.2 Blood Insulin Absorption Compartment Model.....	50
5.2.3 Insulin Kinetics Compartment Model	51
5.2.4 Insulin Signaling Pathway Model	53
5.2.5 Long Short Memory Neural Networks	54
5.3 Model Training and Evaluation	56
5.3.1 Handling Missing Data and Normalizing Features	56
5.3.2 Hyperparameters' tuning.....	56
5.3.3 Performance Assessment of the Model.....	58
5.3.3.1 Binary Classification Threshold Optimization	58
5.3.3.2 Prediction Uncertainty Estimation.....	58
5.3.3.3 External Evaluation of the Model	58
6. Results and Discussion	61
6.1 Discrimination Performance	61
6.2 Comparison of different input space compositions.....	62
6.3 Estimation of Model's Prediction Uncertainty.....	63
6.4 Interpretation of Predictions Using Attention Weights.....	64
6.4.1 Attention Scores Visualization for Characteristic Predictions	67
6.4.1.1 True Positive Prediction	67
6.4.1.2 True Negative Prediction	68
6.4.1.3 False Positive Prediction.....	69
6.4.1.4 False Negative Prediction	70
6.5 Evaluation of the Model on the SMARTDIAB dataset.....	72
6.5.1 Discrimination Performance	72

6.5.2 Estimation of Model's Prediction Uncertainty.....	73
6.5.3 Interpretation of Predictions Using Attention Weights.....	74
6.6 Comparison with the State of the Art.....	76
7. Conclusions and Future Work.....	78
8. References	80

List of Figures

Figure 1: Regulation of blood glucose levels from insulin and glucagon [3].	4
Figure 2: Types of diabetes [8].	6
Figure 3: Chronic diabetes complications [16].	11
Figure 4: Model of a Perceptron [33].	19
Figure 5: Fully connected feedforward neural network (MLP) with two hidden layers [33].	22
Figure 6: A RNN with one input unit, one recurrent hidden unit and one output unit [36].	24
Figure 7: Unraveled form of a RNN [36].	25
Figure 8: The architecture of a LSTM unit [38].	26
Figure 9: The cell state of a LSTM unit [38].	26
Figure 10: The forget gate layer of a LSTM unit [38].	26
Figure 11: The input gate layer of a LSTM unit [38].	27
Figure 12: Update of the cell state in a LSTM unit [38].	27
Figure 13: The output gate of a LSTM unit [38].	28
Figure 14: A visual representation of machine learning types [39].	28
Figure 15: A visual representation of supervised learning [39].	29
Figure 16: A visual representation of unsupervised learning [39].	31
Figure 17: A visual representation of reinforcement learning [39].	32
Figure 18: k-fold cross-validation for k = 10 [43].	34
Figure 19: ROC curve and AUC [46].	36
Figure 20: Diagrammatic representation of the hybrid model for predicting nocturnal hypoglycemic events.	46
Figure 21: Computation of the values for Q_{sto} , Q_{sto1} , Q_{sto2} , Q_{gut} , R_{ag} and k_{empt} over a 120-minutes period, given a carbohydrate intake of 30 grams and a body weight of 85 kilograms, utilizing the Intestinal Glucose Absorption CM.	48
Figure 22: Computation of the values for Q_{sto} , Q_{sto1} , Q_{sto2} , Q_{gut} , R_{ag} and k_{empt} for three different carbohydrate intakes of 25g, 10g and 17g over a 360-minutes period. The calculation was conducted using the Intestinal Glucose Absorption CM and by superimposing the corresponding output curves. ...	49
Figure 23: Computation of the quantities I_{sc1} , I_{sc2} and R_{ai} for three different insulin doses of 2.5, 10 and 6.5 units over a 720-minute period and a body weight of 75 kg, using the Blood Insulin Absorption CM.	51
Figure 24: Computation of insulin plasma concentration $I(t)$ for three different insulin doses of 2.5, 10 and 6.5 units over a 720-minute period and a body weight of 75 kg, using the Insulin Kinetics CM and Blood Insulin Absorption CM (for R_{ait}).	52
Figure 25: Estimation of the percentage of cell surface GLUT4 transporters using the Sedaghat simulation model, following the administration of insulin at a constant concentration of 10^{-5} M, over the time interval from $t = 0$ to $t = 15$ minutes.	53
Figure 26: LSTM-based Model Architecture.	54
Figure 27: Custom attention mechanism structure.	55
Figure 28: Example of StratifiedKFold with k=5 on an imbalanced dataset with two classes [68].	57
Figure 29: Evaluation results of the proposed model.	61
Figure 30: Distribution of Prediction Uncertainties.	63
Figure 31: Uncertainty across the mean predicted probability values.	63

Figure 32: Overlaid histogram of mean probabilities for both correct and incorrect predictions.	64
Figure 33: Visualization of attention weights for the <i>CGM</i> input vectors. The heatmap (left) shows attention distribution across samples and time steps, while the line plot (right) summarizes the total attention at each time step. Timestep “t” refers to the initiation of the sleep period.	65
Figure 34: Visualization of attention weights for the <i>Rag</i> input vectors. The heatmap (left) shows attention distribution across samples and time steps, while the line plot (right) summarizes the total attention at each time step. Timestep “t” refers to the initiation of the sleep period.	65
Figure 35: Visualization of attention weights for the <i>Rai</i> input vectors. The heatmap (left) shows attention distribution across samples and time steps, while the line plot (right) summarizes the total attention at each time step. Timestep “t” refers to the initiation of the sleep period.	66
Figure 36: Visualization of attention weights for the <i>Glut4</i> input vectors. The heatmap (left) shows attention distribution across samples and time steps, while the line plot (right) summarizes the total attention at each time step. Timestep “t” refers to the initiation of the sleep period.	66
Figure 37: Attention of the four input vectors for a TP prediction. Timestep “t” refers to the initiation of the sleep period.	67
Figure 38: Attention of the four input vectors for a TN prediction. Timestep “t” refers to the initiation of the sleep period.	69
Figure 39: Attention of the four input vectors for a FP prediction. Timestep “t” refers to the initiation of the sleep period.	70
Figure 40: Attention of the four input vectors for a FN prediction. Timestep “t” refers to the initiation of the sleep period.	71
Figure 41: Evaluation results of the machine learning model in the SMARTDIAB dataset.	72
Figure 42: Visualization of predictions’ uncertainty in the SMARTDIAB dataset. Top left: Overlaid histogram of mean probabilities for both correct and incorrect predictions. Bottom left: Uncertainty across the mean predicted probability values. Right: Distribution of Prediction Uncertainties	73
Figure 43: Heatmaps of the attention weights of the four input vectors of the SMARTDIAB dataset. Timestep “t” refers to the initiation of the sleep period.	74
Figure 44: Line plots of the averaged attention weights of the four input vectors of the SMARTDIAB dataset. Timestep “t” refers to the initiation of the sleep period.	75

List of Tables

Table 1: Cohort, gender, age group, training data and testing data entries for each contributor of the OhioT1DM dataset [62].	44
Table 2: Mapping of the OhioT1DM dataset.	45
Table 3: Global datasets with the combined data for each patient.	46
Table 4: Mean and variance statistics of the CGM vectors, following the imputation of missing values using the KNNImputer with n_neighbors values of 3, 5, 7, and 10.	56
Table 5: Parameter values checked for hyperparameter optimization.	57
Table 6: Age in years, gender, T1DM duration in years, BMI in kg/m ² and percentage of HbA1c for each patient in SMARTDIAB dataset.	59
Table 7: Descriptive statistics of the SMARTDIAB data.	59
Table 8: Evaluation results of the machine learning models with the different number of input vectors.	62
Table 9: Comparison of model performance metrics in T1D nocturnal hypoglycemia prediction studies. RF: Random Forest, LDA: Linear Discrimination Analysis, SVM: Support Vector Machine, SVR: Support Vector Regression.	76

1. Introduction

1.1. Research Problem

Diabetes Mellitus is a chronic condition characterized by excessively high blood glucose levels, either due to limited secretion of the hormone insulin, reduced action of insulin, or a combination of these two factors. The disease is directly associated with long-term damage, dysfunction, and eventual failure of various vital organs such as the eyes, kidneys, nerves, heart, and blood vessels. According to the World Health Organization, in recent decades, diabetes has taken on epidemic proportions, with cases rising from 108 million in 1980 to 422 million in 2014. There is a globally agreed target to halt the rise in diabetes and obesity by 2025.

To manage this condition, beyond adopting a balanced diet and engaging in physical exercise, regulating blood glucose through regular monitoring and insulin administration is also of crucial importance. However, glycemic control is a demanding and complex process, as glucose levels are influenced by numerous factors, making it ultimately empirical and prone to errors. In recent years, the application of machine learning models has contributed to the development of more accurate methods of glycemic control that take into greater account the individual peculiarities of each case, leading to more efficient and personalized treatment. The further advancement of such artificial intelligence techniques raises hopes for mitigating the complications of the condition and improving the overall quality of life for patients.

One significant challenge in managing diabetes is nocturnal hypoglycemia, which occurs when blood glucose levels drop dangerously low during sleep. This condition can cause symptoms such as sweating, trembling, confusion, and even seizures. If not promptly addressed, nocturnal hypoglycemia can lead to severe health problems, including brain damage and death. The development of continuous glucose monitoring systems and advanced insulin pumps, often integrated with machine learning algorithms, aims to predict and prevent hypoglycemic episodes, offering a safer and more stable management of diabetes, particularly during the night. This technological advancement is essential for reducing the risk of nocturnal hypoglycemia and enhancing the safety and quality of life for individuals with diabetes.

1.2 Objective of the Thesis

The objective of this study is to develop an interpretable model for predicting the risk of nocturnal hypoglycemic episodes using machine learning techniques. This prediction involves estimating the likelihood of a nocturnal hypoglycemic episode occurring within a certain future timeframe during the night, rather than directly forecasting future glucose levels. Furthermore, it is deemed essential to incorporate interpretability techniques into the final system to achieve

a deeper understanding of its operation and result extraction process. Interpretability is increasingly sought after for many modern machine learning models, particularly when applied in healthcare settings to aid medical decision-making. Providing explanations for a model's outputs enhances its reliability, serves as a means of verifying its correct functioning, and facilitates the easier identification of errors that may arise during its application in real-world, non-experimental conditions.

1.3 Structure of the Thesis

This thesis is organized into seven chapters. The initial chapters aim to provide the theoretical foundation for the work by offering a deeper understanding of the nature of diabetes and the mechanisms underlying artificial neural networks. The second part focuses primarily on the experimental process followed. It details the methods developed, the resulting outcomes, and the final conclusions and observations. More specifically:

- Chapter 1 serves as a brief introduction, clarifying the research problem and the motivation behind the study, as well as presenting the overall structure of the thesis.
- Chapter 2 provides a general overview of diabetes, covering the most common types of the disease, its causes and symptoms, and both short-term and long-term complications. It also addresses methods of diagnosis and effective management.
- Chapter 3 introduces key concepts of artificial neural networks, with an emphasis on recurrent neural networks, which are often used in applications like the one under study. It also discusses major training and evaluation methods for machine learning models, many of which are applied in the experimental section.
- Chapter 4 aims to deepen the understanding of the research problem by reviewing recent related literature. It focuses on modern models for predicting glucose levels and hypoglycemic episodes using machine learning.
- Chapter 5 outlines the techniques used for processing the available data and for the development and evaluation of the machine learning model applied in this thesis.
- Chapter 6 presents and discusses the experimental results from the application of the proposed model.
- Chapter 8 summarizes the conclusions of the thesis and offers suggestions for future research directions.

2. Diabetes Mellitus

2.1 Glucose Metabolism and Regulation

2.1.1 Glucose Metabolism: Sources and Pathways

Glucose stands as the foremost monosaccharide, serving as a vital energy source for both plants and animals. In human physiology, glucose is acquired through dietary intake or synthesized via metabolic pathways such as glycogenolysis and gluconeogenesis. When considering food intake, carbohydrates represent the primary external source of glucose. These carbohydrates, polysaccharides found abundantly in various foods, predominantly reside in the starch of plant-derived sources like grains, potatoes, and rice.

Shortly after carbohydrate ingestion, a hydrolysis process commences, breaking down these complex molecules into simpler monosaccharides, primarily glucose. This hydrolysis begins in the stomach and continues within the small intestine. Post-absorption in the small intestine, the resultant glucose enters the hepatic portal vein, promptly absorbed by the liver. Here, a fraction of glucose serves hepatic functions, while the majority enters circulation, destined for distribution across tissues like the brain and muscles.

Facilitated by membrane transporter proteins, known as GLUTs, glucose travels through cell membranes, entering cells' interiors. Once inside, most glucose undergoes immediate utilization for adenosine triphosphate (ATP) synthesis, the organic molecule powering cellular activities. Any excess glucose not utilized for energy production is stored for later use, either as glycogen or fat.

Glycogen, a complex polysaccharide comprised of numerous glucose molecules arranged in a branching pattern, serves as a major reservoir for stored glucose. When required, glycogen is enzymatically broken down into glucose units, a process termed glycogenolysis. Predominantly, glycogen stores reside in the liver, with significant concentrations also found in skeletal muscles.

While carbohydrates primarily fuel glucose production, alternative substrates such as amino acids from proteins and glycerol from triglycerides can also contribute. Particularly in instances of low carbohydrate intake, amino acids can be converted into glucose through gluconeogenesis. Moreover, during fasting periods, glycerol derived from triglycerides can be utilized for glucose synthesis [1, 2].

2.1.2 Regulation of Blood Glucose

The aforementioned points highlight the complexity of glucose homeostasis, given that glucose participates in various reactions involving multiple tissues and organs. It is crucial to maintain blood glucose levels within a narrow range because, under normal conditions, glucose serves as

the primary energy source for most tissues and is the sole nutrient the brain can metabolize. In healthy individuals, blood glucose levels are closely regulated to stay within a specific range, regardless of the meals consumed throughout the day. This regulation is managed by hormones, primarily insulin and glucagon, with additional related hormones identified in recent decades [1, 2].

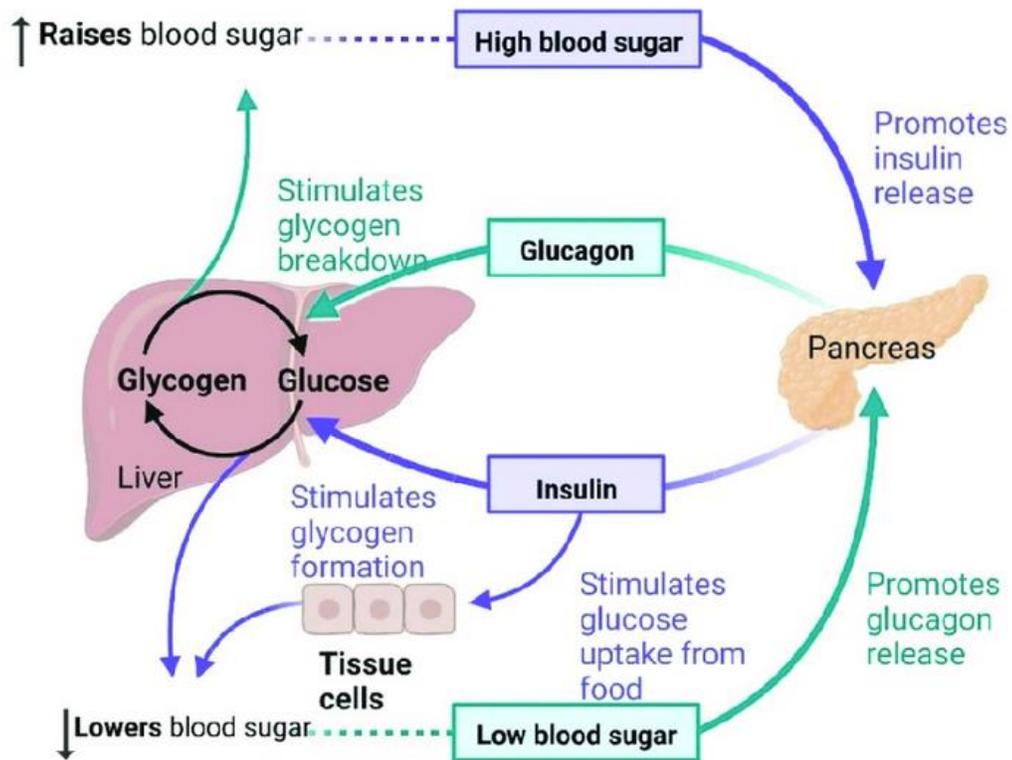


Figure 1: Regulation of blood glucose levels from insulin and glucagon [3].

2.1.2.1 Insulin

Insulin is a typical peptide hormone, consisting of 51 amino acids arranged in two polypeptide chains. It was first isolated by Frederick Banting in 1921. This hormone is produced by the β -cells in the pancreas, one of the four distinct cell types found in the islets of Langerhans. These islets are small cell clusters within the pancreas that are responsible for the secretion of essential hormones. Although they represent less than 2% of the pancreas, β -cells constitute about 75% of the islets. The main function of insulin is to lower blood glucose levels following meals and restore them to their normal state. Insulin also enhances the cellular uptake and storage of glucose, activates enzymes responsible for glycogen, protein, and fat synthesis, and inhibits the production of new glucose via processes such as gluconeogenesis.

The nutrient absorption following a meal increases the plasma concentrations of glucose and amino acids, activating the glucose-sensitive β -cells. Once activated, these cells secrete insulin,

which targets liver cells, brain cells, and cells in adipose and muscle tissues. When insulin binds to the membrane receptors of these target cells, it initiates a cascade of protein phosphorylations inside the cells, leading to the activation of GLUT4 transporters (insulin sensitive). These transporters move from the cytoplasmic vesicles to the cell membrane, allowing substantial amounts of glucose to enter the cells via facilitated diffusion. Consequently, blood glucose levels are reduced and gradually return to their normal range [4, 5].

2.1.2.2 Glucagon

Glucagon, a vital peptide hormone in regulating blood sugar, was first identified in 1922. Comprising 29 amino acids, it is produced by the α -cells in the pancreatic islets of Langerhans. Functioning in opposition to insulin, glucagon is released when blood glucose levels drop, promoting glucose production through glycogenolysis and gluconeogenesis. During the initial 8-12 hours of fasting, glycogen breakdown in the liver serves as the main glucose source. Beyond this period, hepatic gluconeogenesis, utilizing amino acids, becomes a significant glucose production pathway [6].

2.1.2.3 Insulin–Glucagon Balance in Blood Glucose Regulation

The balance between insulin and glucagon is crucial in determining which hormone will dominate and how blood glucose levels will be managed. During fasting periods, the body prevents hypoglycemia by secreting glucagon. When glucagon is predominant, the liver primarily utilizes glycogen to release glucose into the bloodstream. In individuals with normal metabolic function, fasting plasma glucose levels are maintained around 90 mg/dL, with low insulin secretion and relatively high glucagon levels. Following a meal, the increase in plasma glucose inhibits glucagon release and stimulates insulin secretion. This insulin boost facilitates glucose uptake into cells, leading to a rapid return of plasma glucose levels to their fasting state [5].

In cases of diabetes mellitus, the aforementioned metabolic processes are disrupted in one or more ways, resulting in the organism's inability to effectively regulate blood sugar levels. Specifically, in type 1 diabetes mellitus, there is a continuous destruction of pancreatic β -cells by the organism itself, leading to partial or total deficiencies in insulin. As a consequence, postprandial glucose levels increase due to inadequate glucose removal from the blood, poorly regulated hepatic glucose production, and accelerated gastric emptying. In type 2 diabetes mellitus, β -cell function is also impaired, while peripheral tissues exhibit resistance to insulin action, resulting once again in inadequate glycemic control. Finally, disturbances in glucagon production are observed in all types of diabetes mellitus [1, 5]. In the subsequent section, the different types of diabetes mellitus are going to be presented.

2.2 Diabetes Overview and Types

Diabetes Mellitus (DM) is a chronic syndrome with global distribution and serious complications for those affected. According to the International Diabetes Federation, in 2021, the total cases of adults with DM were 537 million, while related deaths reached 6.7 million. It's worth noting that the total number of affected individuals is expected to rise to 643 million by 2030. Three out of four adults with diabetes live in low or middle-income countries, and half of them remain undiagnosed. In all its forms, DM is characterized by disturbances in the metabolism of carbohydrates, fats, and proteins, resulting in elevated blood glucose levels. It is distinguished into Type 1 Diabetes, Type 2 Diabetes, Gestational Diabetes, and other specialized, rarer types [7].

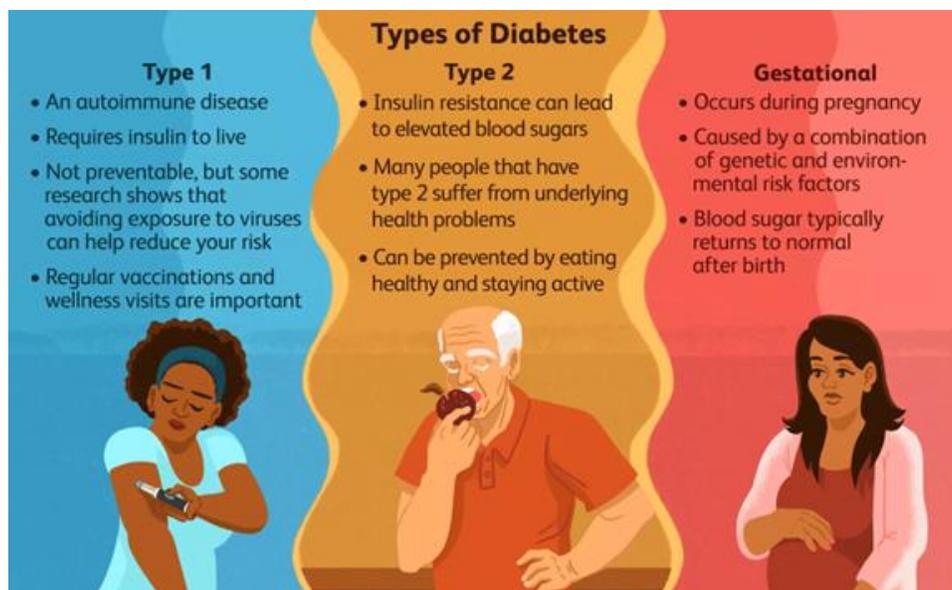


Figure 2: Types of diabetes [8].

2.2.1 Type 1 Diabetes

Type 1 Diabetes accounts for approximately 5%-10% of all diabetes cases. In this type, the patient's own body gradually destroys the pancreatic cells that produce insulin through an autoimmune mechanism. Without this hormone, regulating blood glucose levels is not possible, thus daily insulin administration is necessary from the onset of the disease. Type 1 Diabetes typically appears in children, adolescents, and young adults, but it can also occur in older ages, although this is rarer. The exact causes of its onset are not known, but a family history of autoimmune diseases is a known risk factor. Weight and physical condition are not directly related to its onset, but healthy habits such as regular exercise and adopting a balanced diet are crucial in managing it [9].

2.2.2 Type 2 Diabetes

Type 2 Diabetes constitutes the majority of diabetes cases worldwide (90% - 95%) and is caused by increased resistance of the body to insulin action. Insulin secretion levels may be normal, or even higher in the initial stages of the disease, however, they become insufficient as the disease progresses. In this later phase, insulin production levels decline, necessitating external administration. The most common clinical characteristics of patients with Type 2 Diabetes include obesity, middle or advanced age, a family history of the disease, and gradually increasing hyperglycemia, often with mild or no symptoms. Sedentary lifestyle and unhealthy dietary habits, prevalent in many countries, contribute significantly to the increasing prevalence of Type 2 Diabetes, even among younger age groups. Despite these common clinical features, patients with Type 2 Diabetes exhibit significant heterogeneity: the disease can occur in children, adolescents, or older adults, in overweight or normal-weight individuals, and the severity can range from asymptomatic to life-threatening. This diversity complicates diabetes management for healthcare professionals and patients alike, highlighting the need for further personalization [9].

2.2.3 Gestational Diabetes Mellitus

Gestational Diabetes Mellitus (GDM) is a type of diabetes that occurs only during pregnancy. This form of diabetes affects 3% - 9% of pregnant women and is associated with hormonal changes that occur during pregnancy, especially during the second and third trimesters. Additional factors such as excess weight and heredity also contribute to the risk. Although hyperglycemic symptoms cease after delivery, a significant portion of women who develop GDM later develop Type 2 Diabetes. Therefore, prevention and systematic monitoring are crucial for the years to come. It is also worth noting that if the condition is not effectively treated during pregnancy, there is an increased risk of preterm birth or giving birth to an overweight baby, while the child itself is at risk of developing obesity or Type 2 Diabetes later in life [9].

2.3 Diabetes Symptoms and Complications

2.3.1 Symptoms

One of the primary symptoms of Diabetes is frequent urination, medically known as polyuria. This occurs when elevated levels of glucose in the bloodstream prompt the kidneys to filter and expel the excess sugar, resulting in increased urine production. Alongside polyuria, individuals with diabetes often experience heightened thirst, termed polydipsia, as the body attempts to offset fluid loss caused by frequent urination. Despite a potentially increased appetite and food intake, unexplained weight loss is another common symptom of diabetes. This weight loss can occur due to the body's inability to efficiently utilize glucose for energy, leading to the breakdown

of muscle and fat tissues. Persistent fatigue and weakness are also prevalent among those with diabetes. The inadequate utilization of glucose deprives cells of essential energy, contributing to feelings of tiredness even after ample rest.

Elevated blood sugar levels in diabetes can lead to changes in the shape of the eye's lens, resulting in blurred vision or difficulty focusing. Moreover, these high glucose levels impair circulation and compromise the immune system, slowing down the body's natural healing processes for wounds and injuries. Individuals with diabetes are also at an increased risk of infections, particularly in the urinary tract, skin, and gums, due to the weakened immune response associated with the condition. Despite regular meals, persistent hunger, or polyphagia, may be experienced by individuals with diabetes. This hunger persists as the body struggles to properly utilize glucose for energy, despite sufficient food intake [10].

2.3.2 Acute Complications

Diabetes complications refer to additional health issues that arise as a result of unregulated blood sugar levels in individuals with diabetes. These complications can be categorized into two main types: acute and chronic. Acute complications are rapid-onset health issues commonly seen in diabetes, such as diabetic ketoacidosis, hyperglycemic hyperosmolar state, hypoglycemia [11].

2.3.2.1 Diabetic Ketoacidosis (DKA)

Diabetic ketoacidosis (DKA) is a severe complication of diabetes resulting from critically low insulin levels, predominantly seen in Type 1 Diabetes but also occurring in Type 2 Diabetes. Insufficient insulin leads to a cascade of metabolic disturbances, primarily driven by the release of counter-regulatory hormones like glucagon. This hormonal response triggers the excessive production of free fatty acids (FFAs) from adipose tissue. The liver subsequently converts these FFAs into ketone bodies through ketogenesis. The accumulation of ketone bodies in the bloodstream leads to a condition known as ketonemia, characterized by elevated levels of ketones in the blood. This metabolic state contributes to a decrease in blood pH, leading to acidosis, a hallmark feature of DKA. Concurrently, the lack of insulin results in uncontrolled hyperglycemia, as gluconeogenesis and glycogenolysis proceed unabated, while glucose uptake by peripheral tissues is impaired. The combination of hyperglycemia and ketosis results in osmotic diuresis, characterized by the excessive excretion of glucose and ketones in the urine, leading to dehydration and electrolyte imbalances.

Symptoms of DKA usually develop swiftly and encompass increased urination (polyuria), excessive thirst (polydipsia), loss of weight, weakness, as well as feelings of nausea and vomiting. Left untreated, DKA can progress to severe dehydration, electrolyte abnormalities, hypotension, shock, and even coma or death. Management of DKA involves addressing the underlying causes, including dehydration, acidosis, and hyperglycemia, through fluid resuscitation, insulin therapy,

and correction of electrolyte imbalances. Prevention strategies focus on vigilant monitoring of blood glucose and ketone levels, prompt insulin adjustments during illness, and patient education on recognizing symptoms and seeking timely medical intervention [11, 12].

2.3.2.2 Hyperglycemia Hyperosmolar State (HHS)

Hyperglycemia hyperosmolar state (HHS) is an acute complication seen more frequently in individuals with Type 2 Diabetes, contrasting with DKA, which is more common in Type 1 Diabetes. HHS shares several symptoms with DKA. It carries a significantly higher mortality rate, approximately ten times greater than that observed in DKA. Both DKA and HHS occur due to decreased insulin effectiveness, either from a shortage of insulin secretion (as in DKA) or insufficient insulin action (as in HHS). Extremely high blood glucose levels lead to osmotic diuresis, causing water to be drawn out of cells into the blood, resulting in increased blood osmolarity and dehydration if not promptly replaced. Electrolyte imbalances further exacerbate the condition.

Unlike DKA, HHS does not typically result in significant ketosis and acidosis due to the minimal presence of insulin suppressing counterregulatory hormones and limiting ketone production. Various factors can trigger HHS, including infection, myocardial infarction, trauma, and certain medications. Urgent medical treatment is necessary, usually beginning with fluid volume replacement. Overall, HHS is characterized by hyperglycemia, hyperosmolarity, dehydration, and mild or no ketosis, constituting a medical emergency requiring immediate attention and intervention [11, 12].

2.3.2.3 Hypoglycemia

Hypoglycemia is a frequent and acute complication of diabetes, particularly affecting those who receive insulin externally. It is defined by blood glucose levels dropping below 70mg/dl, presenting with symptoms like discomfort, weakness, sweating, hunger, confusion, irritability, rapid heart rate, and headaches. In serious cases, the patient needs help from others to manage the condition. Critical complications can include cardiac arrhythmias, seizures, coma, and rarely, death. Asymptomatic episodes, known as hypoglycemia unawareness, can also occur. Common causes include excessive doses of antidiabetic drugs, skipping meals, unexpected physical activity, and alcohol consumption while fasting. Treatment involves administering glucose either directly or through food. Preventing hypoglycemia requires constant monitoring and consistent management of blood glucose levels [12].

2.3.2.4 Nocturnal Hypoglycemia

Nocturnal hypoglycemia is a prevalent and perilous complication in patients with diabetes, particularly those with Type 1 Diabetes undergoing insulin therapy. This condition is

characterized by abnormally low blood glucose levels during sleep, often going unnoticed by patients due to the natural suppression of hypoglycemia awareness during sleep. The consequences of nocturnal hypoglycemia are severe, ranging from minor disruptions such as nightmares and morning headaches to major health risks including seizures, cardiac arrhythmias, and the potentially fatal "dead-in-bed" syndrome.

One of the primary dangers of nocturnal hypoglycemia is its tendency to occur without waking the patient. This lack of awareness is particularly concerning as it can lead to prolonged periods of hypoglycemia, exacerbating the risk of severe outcomes. Studies have shown that up to 50% of severe hypoglycemic episodes occur during the night, emphasizing the high vulnerability during sleep. The immediate effects of such episodes can include convulsions, which not only pose a direct physical risk but also contribute to heightened fear and anxiety surrounding hypoglycemia for both patients and their families. Furthermore, nocturnal hypoglycemia has long-term implications on cognitive function. Recurrent episodes can diminish the body's counterregulatory responses to hypoglycemia, leading to impaired awareness and increased risk of future episodes. This condition, known as hypoglycemia unawareness, creates a vicious cycle that complicates diabetes management and endangers patients.

Preventive strategies are crucial in managing nocturnal hypoglycemia. Regular blood glucose monitoring, particularly at bedtime and during the night, can help identify at-risk periods. Adjusting insulin regimens to minimize nocturnal hyperinsulinemia, alongside consuming long-acting carbohydrates before sleep, can stabilize blood glucose levels overnight. The use of continuous glucose monitoring systems has also shown promise in providing real-time alerts and preventing severe nocturnal hypoglycemia.

In summary, nocturnal hypoglycemia poses significant challenges due to its subtle presentation, severe potential outcomes, and the difficulty in detection and prevention. Comprehensive management strategies, including regular monitoring and tailored insulin regimens, are essential to mitigate these risks and improve the safety and quality of life for patients with diabetes [13, 14].

2.3.3 Long-Term Complications

Chronic complications are the leading cause of death among people with diabetes mellitus. However, with diligent management and a healthy lifestyle, these complications can often be prevented or slowed. The discussion here highlights issues affecting the heart and major blood vessels, kidneys, eyes, and peripheral limbs. Additionally, diabetes can lead to a range of other problems, including digestive, urological, sexual, and mental health disorders [15].

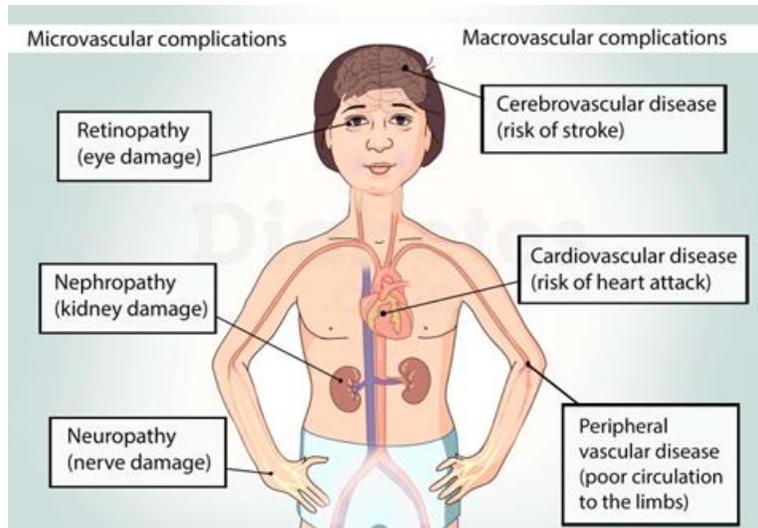


Figure 3: Chronic diabetes complications [16].

2.3.3.1 Diabetic Retinopathy

Diabetic retinopathy is a serious eye condition that can lead to vision loss and blindness among individuals with diabetes. It primarily affects the blood vessels in the retina, the light-sensitive tissue at the back of the eye responsible for detecting light and transmitting signals to the brain via the optic nerve. The underlying mechanism of diabetic retinopathy involves high blood sugar levels damaging the blood vessels in the retina. This damage results in poor blood flow, leading to blocked vessels that leak fluid or bleed. In response, the eye attempts to grow new blood vessels, but these are often abnormal and prone to leaking or bleeding, exacerbating the condition.

In the early stages of diabetic retinopathy, individuals often do not experience any symptoms. Some may notice intermittent changes in their vision, such as difficulty reading or seeing distant objects, but these changes can fluctuate. As the disease progresses, the blood vessels in the retina begin to leak or bleed into the vitreous, the gel-like substance filling the eye. This bleeding can cause the appearance of dark, floating spots or cobweb-like streaks in one's vision. These spots may temporarily clear, but without timely treatment, scarring can occur, further damaging the retina and potentially leading to more severe bleeding.

Diabetic retinopathy can lead to several other serious eye conditions. One such condition is diabetic macular edema (DME), which affects approximately 1 in 15 people with diabetes. DME occurs when blood vessels leak fluid into the macula, the part of the retina responsible for sharp central vision, causing blurry vision. Another complication is neovascular glaucoma, where abnormal blood vessels grow out of the retina and block the eye's drainage system, leading to increased eye pressure and potential vision loss. Additionally, retinal detachment can occur when

scar tissue caused by diabetic retinopathy pulls the retina away from the back of the eye, a condition known as tractional retinal detachment.

Effective management of diabetic retinopathy focuses on controlling blood sugar levels to prevent or slow the progression of the disease. This includes regular physical activity, a healthy diet, and adherence to prescribed insulin or other diabetes medications. Early detection and timely treatment are crucial, as they can prevent significant vision loss and improve the quality of life for individuals with diabetes. Treatments for advanced stages may include laser therapy, vitrectomy, or injections to reduce swelling and inhibit the growth of abnormal blood vessels [15, 17].

2.3.3.2 Diabetic Nephropathy

Diabetic nephropathy, also known as diabetic kidney disease, is a serious complication of diabetes that affects the kidneys' ability to filter waste from the blood. This condition develops in approximately 30-50% of patients with diabetes, particularly those with Type 1 and Type 2 Diabetes. The primary cause is the prolonged high blood glucose levels that damage the tiny blood vessels in the kidneys, impairing their filtering capability.

Early symptoms of diabetic nephropathy are often subtle and include the presence of the protein albumin in the urine, known as albuminuria. As the condition progresses, patients may experience more pronounced symptoms such as edema in the legs and ankles, weight gain, increased frequency of urination, elevated blood pressure, fatigue, and nausea. If left untreated, diabetic nephropathy can lead to kidney failure, necessitating dialysis or a kidney transplant. The pathophysiology of diabetic nephropathy involves complex interactions between hemodynamic changes, metabolic dysregulation, and inflammatory processes.

Controlling blood pressure and blood glucose levels is the main goal for controlling diabetic nephropathy in order to delay the disease's progression. Regular monitoring of kidney function through blood tests and urine analysis is crucial for early detection and management of diabetic nephropathy. Alongside pharmacological treatments, lifestyle modifications such as maintaining a healthy diet, regular physical activity, and avoiding smoking are essential to managing diabetic nephropathy [15, 18].

2.3.3.3 Diabetic Cardiovascular Disease

Diabetic cardiovascular disease (CVD) represents a significant complication for individuals with diabetes, contributing to a high risk of mortality and morbidity. This condition encompasses a range of cardiovascular disorders, including coronary artery disease, heart failure, stroke, and peripheral arterial disease. The pathophysiology of diabetic CVD involves multiple interrelated processes. Chronic hyperglycemia leads to the formation of advanced glycation end-products

(AGEs), which damage the endothelium and promote oxidative stress. This endothelial damage impairs vasodilation and contributes to the development of atherosclerotic plaques. Additionally, diabetes-induced dyslipidemia (characterized by high levels of LDL cholesterol and low levels of HDL cholesterol) exacerbates plaque formation. Hypertension, common in diabetic patients, further strains the cardiovascular system by increasing arterial stiffness and promoting left ventricular hypertrophy.

Diabetic CVD often manifests as chest pain, shortness of breath, fatigue, and in severe cases, heart attack or stroke. Peripheral arterial disease presents with symptoms such as leg pain during walking (claudication) and, in advanced stages, can lead to gangrene and amputation. Heart failure symptoms include persistent coughing or wheezing, swelling in the legs, ankles, and feet, and rapid or irregular heartbeat.

Managing diabetic CVD involves a multifaceted approach focused on controlling blood glucose and blood pressure, as well as addressing lipid abnormalities. Pharmacological treatments include antihypertensive medications, cholesterol management drugs, and antiplatelet agents to prevent clot formation. Lifestyle modifications are equally crucial; these include adopting a healthy diet, engaging in regular physical activity and maintaining a healthy weight [15, 19].

2.3.3.4 Diabetic Neuropathy

Diabetic neuropathy is a common complication of diabetes that affects the nerves due to chronic high blood sugar levels. It encompasses a range of nerve disorders that vary in severity and can significantly impact the quality of life for individuals with diabetes. The symptoms of diabetic neuropathy depend on the type and nerves affected. Peripheral neuropathy, the most common form, typically presents pain, tingling, or numbness in the extremities, especially the feet and legs. Autonomic neuropathy affects involuntary bodily functions, leading to issues such as digestive problems, urinary difficulties, and cardiovascular abnormalities. Proximal neuropathy causes pain in the thighs, hips, or buttocks and can lead to weakness in the legs. Lastly, focal neuropathy results in sudden weakness of one nerve or a group of nerves, causing muscle weakness or pain.

The pathogenesis of diabetic neuropathy involves multiple factors. Chronic hyperglycemia leads to metabolic and vascular changes that damage nerves. High blood sugar causes AGEs to accumulate, oxidative stress, and inflammation, all contributing to nerve damage. Additionally, poor blood flow to nerves due to damaged blood vessels exacerbates this injury, impairing the nerves' ability to transmit signals. Diabetic neuropathy can lead to severe complications if not properly managed. Peripheral neuropathy increases the risk of foot ulcers and infections, potentially leading to amputations due to poor healing. Autonomic neuropathy can cause a range of issues, from gastrointestinal symptoms like gastroparesis to cardiovascular problems like

orthostatic hypotension, which can increase the risk of falls and fractures. Additionally, neuropathy-related pain can significantly impact sleep and daily activities, reducing overall quality of life.

Managing diabetic neuropathy focuses on controlling blood sugar levels to prevent further nerve damage. This can be achieved through lifestyle modifications such as a healthy diet, regular exercise, and adherence to prescribed diabetes medications. Blood pressure and cholesterol management are also crucial, as hypertension and dyslipidemia can worsen neuropathy [15, 20].

2.4 Diabetes Diagnosis

The appearance of symptoms or certain complications often leads to the diagnosis of diabetes mellitus in an individual. Nonetheless, confirming this diagnosis requires a diagnostic test. These clinical diagnostic assessments vary and should always be performed by a specialist. The primary tests used to diagnose diabetes in practice are outlined below. It is important to note that, in most cases, a definitive diagnosis requires the potential patient to undergo multiple tests, spaced a few days apart [21, 22].

2.4.1 Hemoglobin A1C (HbA1c) Test

The HbA1c test measures the percentage of glycated hemoglobin in venous blood. Hemoglobin, found in red blood cells, is responsible for carrying oxygen to tissues, and when it bonds chemically with glucose, it forms glycated hemoglobin (HbA1c). The HbA1c level reflects average blood sugar levels over the past 120 days, thus it can be used to determine the average blood glucose level for the previous three months. An HbA1c level of 6.5% or higher indicates diabetes mellitus, while levels between 5.7% and 6.4% signify prediabetes, a condition that increases the risk of developing diabetes [21, 22].

2.4.2 Fasting Plasma Glucose (FPG) Test

The FPG test measures blood glucose levels after an individual has fasted for at least 8 hours. This test is straightforward, cost-effective, and precise for diagnosing the metabolic imbalance linked to diabetes mellitus. Normal glucose levels are below 100 mg/dl, levels from 100 mg/dl to 125 mg/dl indicate prediabetes, and levels above 126 mg/dl are diagnostic of diabetes mellitus [21, 22].

2.4.3 Glucose Tolerance Test (GTT)

This test measures plasma glucose levels before and after consuming a glucose solution, aiming to evaluate glucose tolerance in an individual. Like the FPG test, it requires an 8-hour fast prior to the assessment. According to the American Diabetes Association, postprandial plasma glucose

levels below 140mg/dl are considered normal, levels between 140mg/dl and 199mg/dl indicate prediabetes, and levels exceeding 200mg/dl suggest diabetes mellitus [21, 22].

2.5 Diabetes Management

Managing diabetes mellitus is a multifaceted and continuous endeavor, demanding vigilant observation of the illness's advancement from both the patient and a healthcare professional. As mentioned earlier, the disease's progression is daily influenced by numerous factors, including the patient's lifestyle, dietary habits, medication adherence, psychological stress levels, comorbidities, and environmental factors. Nonetheless, consistent and informed diabetes management can uphold a decent quality of life and mitigate the emergence of acute and chronic complications [23, 24].

2.5.1 Blood Glucose Monitoring

Regularly monitoring blood glucose levels is crucial for individuals with diabetes and is an essential aspect of their daily lives. This practice offers numerous advantages, including better control over meal timing and medication administration, as well as reducing the risk of severe blood sugar fluctuations. Analyzing blood glucose data over time provides valuable insights into disease progression and helps healthcare providers make informed decisions about necessary lifestyle adjustments.

The traditional method of glucose monitoring involves obtaining a small blood sample, typically from the patient's fingers, multiple times throughout the day using specialized devices. Monitoring frequency varies depending on diabetes type, individual factors, and treatment regimen. For instance, those with Type 1 Diabetes may need to check their levels 4-10 times daily, emphasizing the importance of measurements before meals and bedtime. Normal blood glucose levels for someone with diabetes are typically below 180mg/dl after meals and within the range of 80mg/dl – 130mg/dl during fasting periods [25].

Recent advancements have introduced Continuous Glucose Monitoring (CGM) devices, which automatically record glucose levels, enhancing the quality of glycemic control, as supported by recent research. These devices utilize sensors typically attached to the arm or abdomen, measuring glucose levels in the interstitial fluid between cells every few minutes. Although these readings may differ slightly from traditional blood glucose measurements, they accurately represent blood glucose levels and their fluctuations. These sensors usually require replacement every few days, are largely waterproof, and often do not need calibration. The effectiveness of these CGM devices is primarily assessed using the Mean Absolute Relative Difference (MARD) metric, with devices showing MARD values close to 10% considered reliable and suitable for practical use.

One significant benefit of CGM technologies is their capacity to continuously track users' glucose levels overnight, a capability lacking in traditional monitoring methods. Moreover, many sensors offer audible alerts to patients in case their glucose levels surpass preset high or low thresholds. Users can customize these thresholds based on guidance from their healthcare provider, enhancing the individualization of their treatment approach. Additionally, storing measurements over time enables data analysis to derive valuable insights and facilitates the development of prediction systems, similar to the one under investigation in this study [26].

CGM sensors are classified into two main categories: rtCGM (real-time CGM) and isCGM (intermittently scanned CGM) based on whether the recorded values are shown in real-time on the device or only upon user request, possibly on a mobile phone. Examples of rtCGM models include Dexcom G6, Dexcom G5, Medtrum Touch Care Nano, and Medtronic Guardian 3/4 sensors, while FreeStyle Libre and FreeStyle Libre 2 models exemplify isCGM devices. Unlike rtCGM models, isCGM devices do not display real-time data unless the user scans the sensor, usually using a mobile phone [27].

2.5.2 Insulin and Antidiabetic Drugs

Administering insulin externally is vital for managing diabetes mellitus. For Type 1 Diabetes, external insulin is essential from the start, due to the complete absence of insulin. For Type 2 Diabetes, initial treatment usually involves oral antidiabetic drugs, with insulin being introduced only when these drugs no longer suffice and the body's insulin production decreases. One way to administer insulin is through Multiple Daily Injection (MDI) therapy, which involves injecting insulin into the abdomen, legs, or arms several times daily using a syringe or an insulin pen. The pen's needles are very small, making injections almost painless. The specific treatment plan is tailored by a doctor according to the type of diabetes and the patient's daily routine. Typically, long-acting insulin is taken before sleeping, and rapid-acting insulin is taken before each meal.[28]

Alternatively, an insulin pump can be used for diabetes management. This device provides a continuous and gradual supply of rapid-acting insulin in small doses throughout the day, mimicking the function of a healthy pancreas. The user can program the basal rate of insulin delivery. The insulin pump system includes a small subcutaneous catheter connected to the pump via a tube, or in the case of "patch pumps," the catheter and insulin reservoir are attached directly to the skin without tubing, programmed wirelessly via a remote device.

For managing meals or high glucose levels, insulin pumps allow for additional doses, known as insulin boluses. These pumps can calculate the necessary dose based on current glucose levels, the amount of carbohydrates to be consumed, and the insulin already present in the body from previous doses. The patient can then use this calculated dose or adjust it as needed. When paired

with a CGM device, some pumps can automatically stop insulin delivery to prevent hypoglycemia. Advanced systems aim to automate insulin delivery based on rtCGM data, aspiring to create an “artificial pancreas” [28, 29].

2.5.3 Nutrition and Exercise

Proper diet and physical activity are essential for managing diabetes. A well-planned diet helps in adjusting medication effectively to maintain optimal blood sugar levels after meals, thereby improving overall glycemic control. Modern guidelines suggest that people with diabetes can enjoy a variety of foods as part of a healthy diet, which is also recommended for everyone. The American Diabetes Association advises that people with diabetes should eat fruits and vegetables daily, consume legumes and cereals regularly, and include fish in their diet weekly. Red meat should be limited to once a week. It's also crucial to avoid simple carbohydrates like sugar, white bread, and pasta, and instead, consume high-fiber carbohydrates. These foods slow digestion and help maintain a sense of fullness. For precise insulin or antidiabetic medication dosing, patients are often required to count the carbohydrate content of their meals beforehand [30].

Engaging in physical activity is highly beneficial for managing diabetes. Weight loss is particularly effective for overweight or obese individuals with diabetes, aiding in better disease management and even prevention in Type 2 Diabetes. Along with a balanced diet, it is recommended that people perform moderate to intense aerobic exercise for at least 30 minutes daily, either continuously or in intervals, at least five days a week. Those with diabetes should minimize sedentary habits and take regular breaks if prolonged sitting is necessary for work. Exercise should be approached carefully, especially for those with Type 1 Diabetes, and avoided if blood glucose levels are too high or too low. Planned exercise requires reducing the insulin dose beforehand, and consuming carbohydrates during extended exercise sessions is advisable [31].

3. Artificial Neural Networks

3.1 Emulating Brain Function for Machine Learning

The brain's functioning has inspired the creation of artificial neural networks (ANNs), a vital part of machine learning. ANNs are made up of small computational units called artificial neurons. Essentially, ANNs are designed to replicate how the brain performs tasks. They operate as distributed systems with high parallelism, capable of storing and using knowledge gained empirically. ANNs emulate the brain by learning from the environment, storing knowledge in weights that define connections between neurons (synaptic weights). Training an ANN involves adjusting these weights until they achieve the desired outcomes, using specific algorithms.

ANNs derive their computational power from their parallel and distributed operations, and their ability to generalize allows them to provide reasonable solutions to problems they haven't encountered before. They are used extensively in diverse applications such as facial recognition, predicting data like stock prices or weather, filtering spam emails, aiding medical decisions, detecting fake news on social media, and automating drug discovery. ANNs are integral in fields like computer vision, natural language processing, and speech recognition, highlighting their effectiveness and broad integration into everyday life [32].

3.2 Perceptron

3.2.1 Origin and Function

The perceptron, pivotal in the historical evolution of neural networks, marked the first algorithmically defined model. Invented by psychologist Rosenblatt, its introduction spurred interdisciplinary research across engineering, physics, and mathematics throughout the 1960s and 1970s. Functionally, the perceptron serves as a basic neural network for classifying linearly separable patterns, employing a single neuron with adjustable synaptic weights and bias. Rosenblatt's algorithm for adjusting these parameters demonstrated the perceptron's ability to converge when training patterns are linearly separable, positioning a hyperplane decision surface between classes. This convergence proof, known as the perceptron convergence theorem, validated its utility. Initially limited to binary classification, expanding the perceptron's output layer allows for multiclass classification by incorporating additional neurons [33].

3.2.2 Fundamentals of the Perceptron

The perceptron serves as the central building block of an ANNs, laying the foundation for the development of more sophisticated and expansive systems that are explored thereafter. Below is the model that outlines the structure of a perceptron:

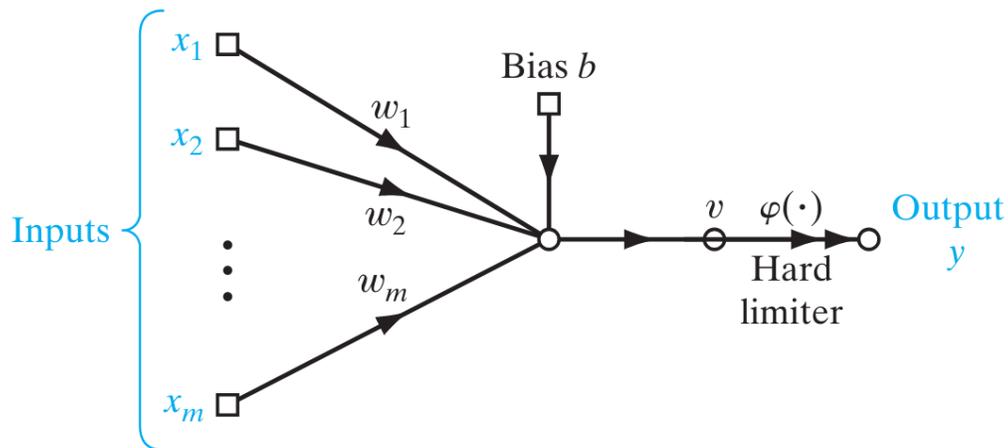


Figure 4: Model of a Perceptron [33].

The perceptron described above is characterized by m connections (or synapses), each associated with a specific weight w_i . Every input signal undergoes multiplication by its respective weight w_i , and these weighted inputs are subsequently summed to produce the value v :

$$v = \sum_{i=1}^m w_i \cdot x_i \quad (1)$$

Additionally, the model incorporates an externally adjustable variable known as bias b , which adjusts the output v , to increase or decrease based on whether it is positive or negative. Lastly, the activation function $\varphi(\cdot)$ restricts the perceptron's output, thereby determining the final result:

$$y = \varphi(v + b) \quad (2)$$

The ANNs composed of a single perceptron are called Perceptrons. Perceptrons are capable of classifying patterns that can be separated by a straight line into two categories. In binary classification, if the sum $v = \sum_{i=1}^m w_i \cdot x_i$ is positive, the perceptron output is +1, if negative, it outputs -1. Therefore, based on whether the perceptron outputs +1 or -1, the point defined by inputs x_1, x_2, \dots, x_m is classified into class $C1$ or $C2$, respectively.

The synaptic weights w_1, w_2, \dots, w_m are adjusted iteratively until they satisfy the condition $w^T \cdot x > 0$ for all input vectors x belonging to class $C1$, and $w^T \cdot x < 0$ for all input vectors x belonging to class $C2$. The weights are updated according to the following rule:

$$w(n+1) = w(n) + \eta \cdot (d(n) - y(n)) \cdot x(n) \quad (3)$$

Here, n represents the current iteration number, η (where $\eta \in (0, 1)$) denotes the learning rate, $y(n) = \varphi(w^T(n) \cdot x(n))$ indicates the perceptron's output for the iteration n , and $d(n)$ represents the desired output, specifically:

$$d(n) = \begin{cases} +1, & x(n) \in C1 \\ -1, & x(n) \in C2 \end{cases} \quad (4)$$

It has been demonstrated that when classes $C1$ and $C2$ are linearly separable by a hyperplane, the described process converges. This convergence enables the computation of the sought-after vector w after a finite number of iterations. [33]

The next section proceeds to outline several frequently utilized types of activation functions. The activation function $\varphi(\cdot)$ plays a critical role in determining the neuron's ultimate output.

3.2.3 Activation Functions

3.2.3.1 Threshold Function

The threshold function (also known as binary step function) is the most basic type of activation function. According to this function, if the input is negative, the neuron does not activate and outputs zero [33, 34].

$$\varphi(v) = \begin{cases} 1, & v \geq 0 \\ 0, & v < 0 \end{cases} \quad (5)$$

3.2.3.2 Linear Function

When the linear function is used as an activation function, a neuron's output varies directly with its input. However, similar to the threshold function, the linear function is not commonly employed because of its straightforward nature [33, 34].

$$\varphi(v) = v \quad (6)$$

3.2.3.3 Sigmoid Function

The sigmoid function is widely used as an activation function because of its non-linear nature and ability to be differentiated. It outputs values within the range of (0, 1), and its curve resembles the shape of the Latin letter S [33, 34].

$$\varphi(v) = \frac{1}{1 + e^{-v}} \quad (7)$$

3.2.3.4 Hyperbolic Tangent (tanh) Function

The tanh function is both continuous and differentiable, producing outputs that span from -1 to 1. Its symmetric properties around the origin often make it a preferred choice compared to the sigmoid function [33, 34].

$$\varphi(v) = \tanh(v) = \frac{2}{1 + e^{-2v}} - 1 \quad (8)$$

3.2.3.5 Rectified Linear Unit (ReLU) function

Using the ReLU function, a neuron becomes active only when its input is positive, leading to a linear output. One of the major benefits of ReLU is its ability to activate neurons in an artificial neural network asynchronously. It's important to highlight that a refined version of ReLU, known as Leaky ReLU, sets a value of $0.01v$ for $v \leq 0$ and v for $v > 0$. This adjustment effectively prevents the issue of too many neurons becoming inactive [33, 34].

$$\varphi(v) = \max(0, v) \quad (9)$$

3.2.3.6 Softmax Function

The softmax activation function is formed by integrating several sigmoid functions and is ideal for multi-class classification tasks. It is often used as the last activation function of a neural network to normalize the output of a network to a probability distribution over predicted output class. Considering that each sigmoid function's output falls within the range (0, 1), it represents the probability that the input belongs to its associated category [33, 34].

$$\varphi(v)_i = \frac{e^{v_i}}{\sum_{k=1}^k e^{v_k}} \quad (10)$$

Here, k represents the total number of classes to classify, and i ranges inclusively from 1 to k .

3.3 Multilayer Perceptrons

To solve more complex issues, particularly those involving the classification of non-linearly separable classes, layers of perceptrons, known as Multilayer Perceptrons (MLPs), are used. These models feature intermediate (hidden) layers of neurons, which are not directly visible from the network's input and output. Typically, each neuron in a layer receives inputs from the outputs of neurons in the previous layer, and its output serves as the input for neurons in the next layer. Because of this structure, MLPs are termed fully connected neural networks. As information travels from input to output, these networks are also referred to as feedforward networks. Additionally, MLPs employ non-linear activation functions to represent both linear and non-linear relationships between inputs and outputs. It has been demonstrated that a simple multilayer feedforward network can approximate any mathematical function to the desired accuracy if it has a sufficient number of hidden layers [33, 35].

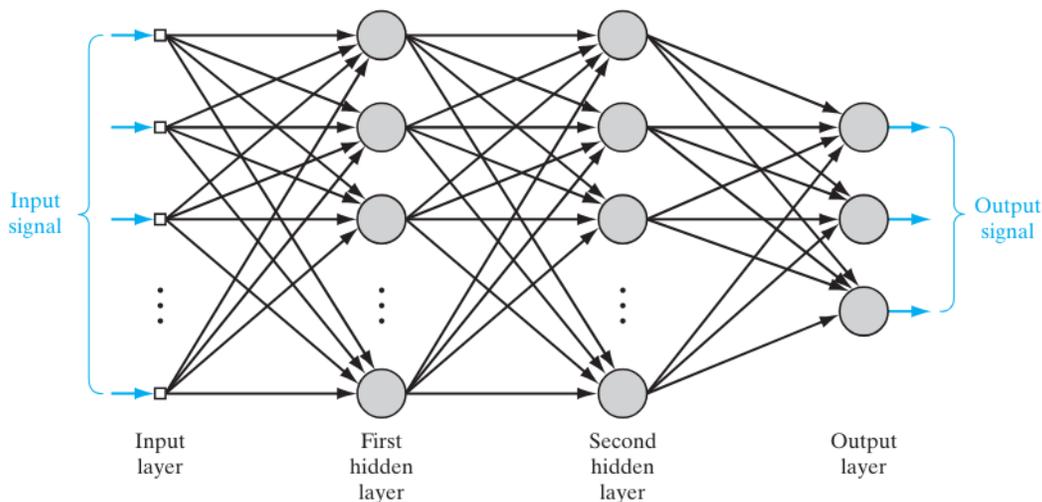


Figure 5: Fully connected feedforward neural network (MLP) with two hidden layers [33].

The distributed non-linearity and high connectivity of MLPs make training them significantly more complex than training simple perceptrons. The most frequently used training method for these networks is the backpropagation algorithm. This algorithm involves two phases: the forward phase and the backward phase. In the forward phase, the synaptic weights are kept constant while information flows through the layers from the input to the output. The network's output is determined by the activation function and the initial input. In the backward phase, an error signal, defined as the difference between the calculated output and the actual output, is propagated back through the network from the output to the input. During this phase, the weights are adjusted layer by layer to minimize this error. Training with the backpropagation algorithm requires knowing the desired output $d(n)$ for each input $x(n)$ beforehand. Specifically, the aim is to calculate the values $\Delta w_{ji}(n)$, which are proportional to the partial derivatives

$\partial E(n)/\partial w_{ij}(n)$, for each input $x(n)$. These derivatives show the necessary adjustments to the synaptic weights $w_{ij}(n)$ to minimize $E(n)$, the total error of the network [33, 35].

We define:

$$E(n) = \frac{1}{2} \sum_{j \in \mathcal{C}} e_j^2(n) \quad (11)$$

where $e_j(n) = d_j(n) - y_j(n)$ represents the error signal for neuron j , and \mathcal{C} denotes the set containing all the neurons in the output layer. To calculate each quantity $\partial E(n)/\partial w_{ij}(n)$, the chain rule can be used:

$$\frac{\partial E(n)}{\partial w_{ij}(n)} = \frac{\partial E(n)}{\partial e_j(n)} \cdot \frac{\partial e_j(n)}{\partial y_j(n)} \cdot \frac{\partial y_j(n)}{\partial v_j(n)} \cdot \frac{\partial v_j(n)}{\partial w_{ij}(n)} \quad (12)$$

Considering that $v_j(n) = \sum_{i=0}^m w_{ij}(n) \cdot x_i(n)$ and $y_j(n) = \varphi_j(v_j(n))$, and since the activation function $\varphi(\cdot)$ is differentiable, we derive the following relationship:

$$\frac{\partial E(n)}{\partial w_{ij}(n)} = -e_j(n) \cdot \varphi_j'(v_j(n)) \cdot y_j(n) \quad (13)$$

The adjustment $\Delta w_{ji}(n)$ made to the weight $w_{ji}(n)$ is governed by the delta rule:

$$\Delta w_{ij}(n) = -\eta \cdot \frac{\partial E(n)}{\partial w_{ij}(n)} \quad (14)$$

where η represents the learning rate in the backpropagation algorithm. The negative sign in the above equation is due to the need to adjust the synaptic weight in a way that decreases $E(n)$, as per the gradient descent method. Therefore, we obtain:

$$\Delta w_{ij}(n) = \eta \cdot \delta_j(n) \cdot y_j(n) \quad (15)$$

where $\delta_j(n) = -e_j \cdot \varphi_j'(v_j(n))$.

Thus, the remaining task is to compute $e_j(n)$, and we need to consider different cases. In the simplest scenario, where j is an output layer neuron, the value of $e_j(n)$ is known and given by $e_j(n) = d_j(n) - y_j(n)$. However, in the more complex scenario where j is a hidden layer neuron, the following recursive relationship is valid:

$$\delta_j(n) = \varphi_j'(v_j(n)) \cdot \sum_{k \in \mathcal{C}} \delta_k(n) \cdot w_{kj}(n) \quad (16)$$

where C represents the set of neurons in the next layer that are connected to the hidden neuron j [33].

3.4 Recurrent Neural Networks

Recurrent Neural Networks (RNNs) are a class of artificial neural networks designed for processing sequential data by maintaining a hidden state that captures information from previous inputs. Unlike traditional feedforward networks, RNNs have connections that form directed cycles, allowing them to process data sequences where context and order are crucial, such as in language modeling, speech recognition, and time series prediction. The core of an RNN involves recursively feeding back the outputs from previous time steps as inputs for the current time step, using a hidden state that evolves over time based on new input data. This recurrent structure enables RNNs to capture temporal dependencies, making them particularly effective for sequence learning tasks.

RNNs have found success in various applications due to their ability to model sequential data effectively. In natural language processing (NLP), RNNs are used for tasks such as language modeling, machine translation, and text generation. In speech recognition, RNNs transcribe spoken language into text, leveraging their ability to handle variable-length inputs and maintain temporal context. They are also used in time series prediction in financial markets and weather forecasting, where their recurrent structure allows them to learn and extrapolate patterns over time.

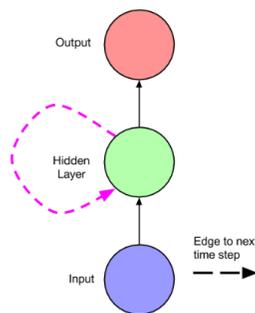


Figure 6: A RNN with one input unit, one recurrent hidden unit and one output unit [36].

The fundamental structure of an RNN consists of an input layer, a hidden layer, and an output layer. At each time step t , the hidden state h_t is updated based on the current input x_t and the previous hidden state h_{t-1} . This process can be mathematically described by the equation:

$$h_t = f(W_{hh} \cdot h_{t-1} + W_{xh} \cdot x_t + b_h) \quad (17)$$

where W_{hh} and W_{xh} are weight matrices, b_h is a bias vector, and f is a non-linear activation function, typically a tanh or ReLU. The hidden state then influences the output y_t , which is computed as:

$$y_t = g(W_{hy} \cdot h_t + b_y) \quad (18)$$

where W_{hy} is the output weight matrix and b_y is the output bias vector and g is often a softmax function for classification tasks.

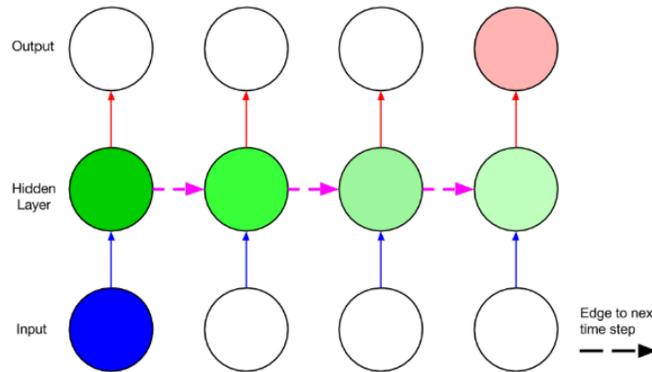


Figure 7: Unraveled form of a RNN [36].

However, training RNNs can be challenging due to issues such as vanishing and exploding gradients, which occur during the backpropagation through time process. These problems can hinder the network's ability to learn long-range dependencies, limiting its effectiveness on longer sequences. To address these issues, more advanced variants of RNNs have been developed, such as Long Short-Term Memory networks and Gated Recurrent Units [36].

3.4.1 Long – Short Term Memory Networks

In 1997, Long Short-Term Memory (LSTM) networks were introduced to handle long-term dependencies and the vanishing gradients problem. While LSTMs retain the chain structure typical of all recurrent neural networks, they offer a more intricate method of building the recurrent units (cells) that form this chain [37].

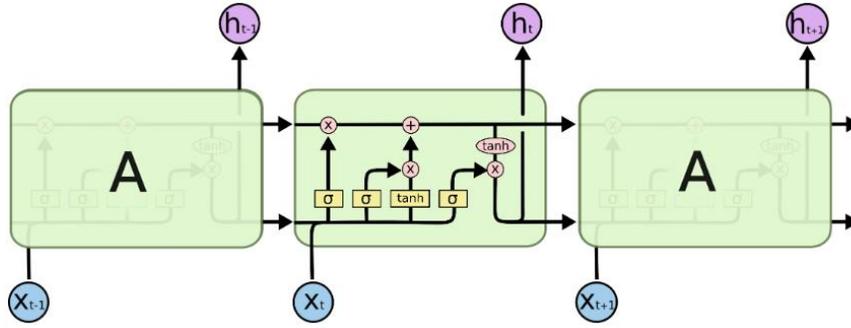


Figure 8: The architecture of a LSTM unit [38].

A core aspect of LSTM networks is the cell state, denoted as C_t for a given cell t . The cell state enables the preservation and direct transfer of past information, requiring only minimal linear adjustments.

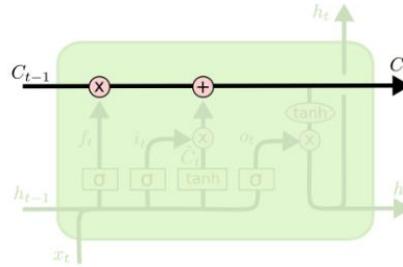


Figure 9: The cell state of a LSTM unit [38].

The C_t in an LSTM network is influenced by the gates f_t (forget gate), i_t (input gate), and o_t (output gate). The initial step in an LSTM involves determining which information will be discarded from the cell state. This decision is facilitated by a sigmoid layer known as the "forget gate layer" (f_t). This layer examines h_{t-1} and x_t , and produces an output ranging between 0 and 1 for each number in the previous cell state C_{t-1} [38].

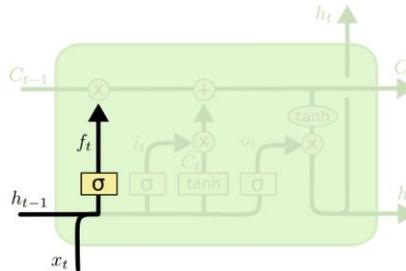


Figure 10: The forget gate layer of a LSTM unit [38].

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \quad (19)$$

where σ is the sigmoid function, W_f is a weight matrix and b_f is a bias vector.

The subsequent step involves deciding what new information will be stored in the cell state. This process comprises two parts. Initially, a sigmoid layer termed the "input gate layer" (i_t) determines which values will be updated. Following this, a tanh layer generates a vector of new candidate values, \tilde{C}_t , that could potentially be added to the state. In the next step, these two elements will be combined to produce an update to the cell state [38].

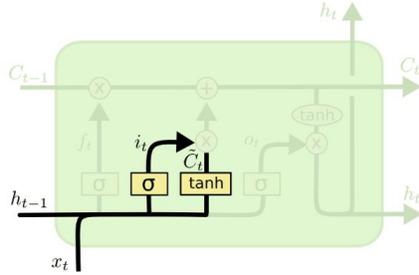


Figure 11: The input gate layer of a LSTM unit [38].

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \quad (20)$$

$$\tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \quad (21)$$

where σ is the sigmoid function, W_i and W_C are weight matrices and b_i and b_C are bias vectors.

Next, the task is to update the old cell state C_{t-1} to the new cell state C_t . The old state is multiplied by f_t , thereby forgetting the information that was previously determined to be unnecessary. Subsequently, $i_t \cdot \tilde{C}_t$ is added. This represents the new candidate values, scaled according to the extent to which each state value is updated [38].

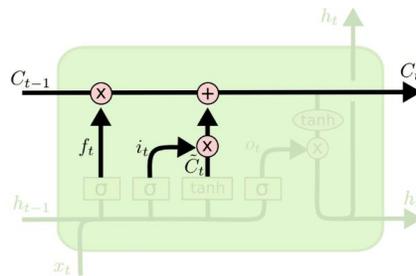


Figure 12: Update of the cell state in a LSTM unit [38].

$$C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t \quad (22)$$

Finally, the output needs to be determined. This output will be based on the cell state but will represent a filtered version. Initially, a sigmoid layer is used to decide which parts of the cell state will be output. Then, the cell state is passed through a tanh function (to constrain the values

between -1 and 1) and multiplied by the output of the sigmoid gate, ensuring that only the selected parts are output [38].

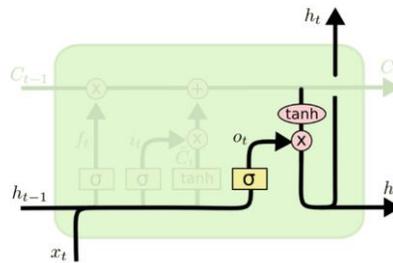


Figure 13: The output gate of a LSTM unit [38].

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \quad (23)$$

$$h_t = o_t \cdot \tanh(C_t) \quad (24)$$

where W_o is a weight matrix and b_o is a bias vector.

3.5 Machine Learning Types

Generally, the learning processes of neural networks can be categorized into two main types: learning with a teacher (supervised learning) and learning without a teacher. Learning without a teacher can be further divided into unsupervised learning and reinforcement learning [33].

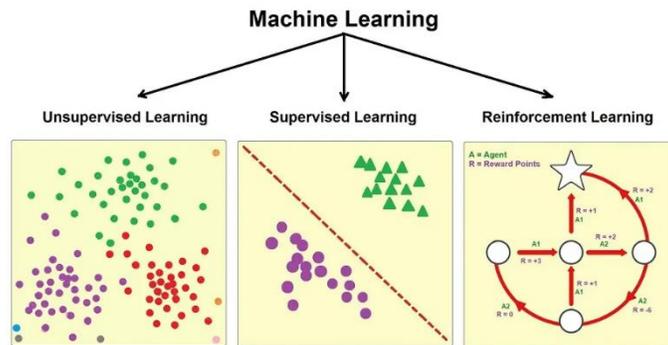


Figure 14: A visual representation of machine learning types [39].

3.5.1 Supervised learning

Supervised learning involves training a model with labeled data. This form of learning is analogous to a teacher-student dynamic, where the teacher (the training data) provides the correct answers, and the student (the machine learning model) learns from these examples to make accurate predictions or decisions. In supervised learning, we start with a set of input-output pairs, known as the training set. Each pair consists of an input vector and the corresponding

desired output. The learning process involves adjusting the model's parameters so that its predictions closely match the desired outputs for the training examples. This adjustment is driven by an error signal, which quantifies the difference between the predicted and actual outputs. The model iteratively updates its parameters to minimize this error, gradually improving its performance [33, 39].

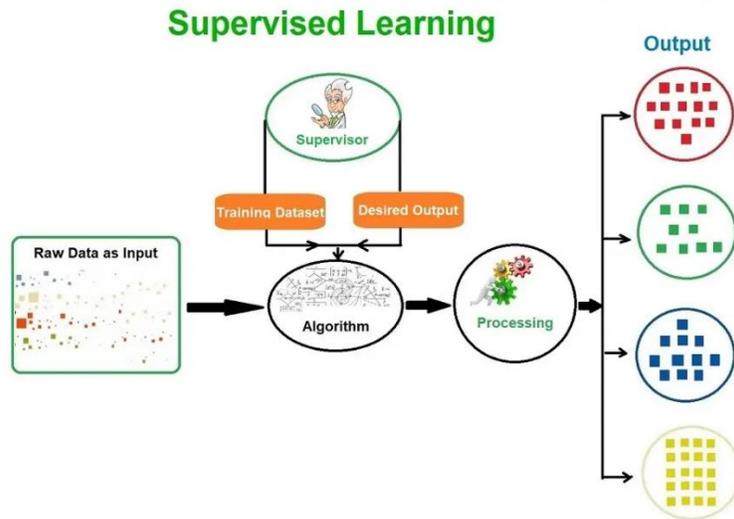


Figure 15: A visual representation of supervised learning [39].

The error-correction mechanism is at the heart of supervised learning. It works by iteratively adjusting the model parameters based on the error signal until the model's output aligns closely with the desired output. This process can be visualized as navigating an error surface, where the goal is to find the point of minimum error. Techniques such as gradient descent are commonly used to navigate this error surface efficiently, ensuring the model converges to an optimal or near-optimal solution.

Supervised learning is frequently applied in classification tasks, where new data inputs are assigned to predefined, distinct categories based on patterns learned by the model during training. Classification can be binary, with two possible output classes, or multi-class, with inputs categorized into one of several classes. In scenarios where inputs can belong to multiple categories, it is known as multi-label classification. Common classifiers used in these tasks include ANNs, decision trees, Support Vector Machines (SVMs), as well as algorithms like Naïve Bayes and k-Nearest Neighbors (k-NN).

In cases where the network's desired output is a continuous numerical value, regression, a statistical method, is frequently used for prediction. This approach aims to understand the relationship between multiple independent variables and dependent variables. In machine learning, the independent variables are the input data features, and they influence the network's final output, which acts as the dependent variable [33, 39].

The effectiveness of supervised learning hinges on the quality and quantity of the labeled training data. Additionally, the model's performance can be impacted by the presence of noise in the data or the complexity of the underlying patterns. Overfitting, where the model learns the training data too well but fails to generalize to new data, is a common challenge that must be addressed through techniques like cross-validation and regularization.

Supervised learning has a wide range of applications across various domains. In healthcare, it can be used to predict patient outcomes based on medical history and diagnostic data. In finance, it can help detect fraudulent transactions. In marketing, it can improve customer segmentation and targeted advertising. The versatility and effectiveness of supervised learning make it a cornerstone of modern machine learning practices [33].

3.5.2 Unsupervised Learning

Unsupervised learning is a type of machine learning that operates without the supervision of a teacher. In this learning paradigm, the system is not provided with labeled inputs and corresponding outputs. Instead, it must find structure and patterns from the input data on its own. This form of learning is particularly useful for discovering hidden patterns or intrinsic structures within the data, making it valuable for tasks such as clustering, association, and dimensionality reduction. Unsupervised learning can be understood from two perspectives: bottom-up and top-down approaches.

The bottom-up approach is inspired by principles of self-organization such as self-amplification, competition, and cooperation. These principles help in forming a model of the learning process where local interactions among elements lead to the emergence of global patterns. Examples of unsupervised neural networks employing this approach include Hebbian learning algorithms, where the learning rule adjusts the weights based on the correlation of inputs.

In contrast, the top-down view involves tuning the adjustable parameters of the model analytically. Given a set of unlabeled examples, the aim is to minimize a cost function subject to constraints imposed on the learning process. This view leverages analytical tools from statistical learning theory. A prominent example of this approach is the kernel Principal Component Analysis (PCA), which is used for extracting higher-order statistics from the input data [33, 39].

Unsupervised learning is extensively applied in several areas, including clustering, association, and dimensionality reduction. Clustering involves grouping similar data points, which is crucial for tasks such as customer segmentation, image compression, and bioinformatics. Association involves finding significant relationships between variables in large datasets, often used in market basket analysis. Dimensionality reduction reduces the number of random variables, aiding in model simplification and data visualization. Techniques like PCA are used to convert

high-dimensional data into a lower-dimensional format while retaining as much information as possible [33].

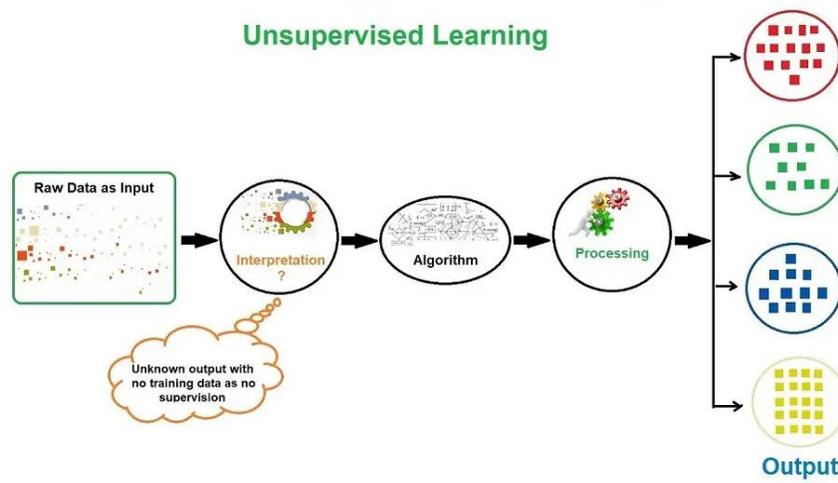


Figure 16: A visual representation of unsupervised learning [39].

3.5.3 Reinforcement Learning

Reinforcement learning is a type of machine learning distinct from supervised and unsupervised learning, focused on how agents should take actions in an environment to maximize some notion of cumulative reward. This method of learning is particularly useful in situations where making decisions is a continuous process, and the agent must learn from the outcomes of its actions without the benefit of a teacher providing explicit feedback at every step.

In reinforcement learning, an agent interacts with its environment through a series of actions and observations. Each action taken by the agent can lead to different states in the environment, which in turn provide feedback in the form of rewards or penalties. The goal of the agent is to learn a policy that maximizes the cumulative reward over time [33, 39].

Reinforcement learning has wide applications, including robotics, game playing, and autonomous systems, where it enables agents to learn complex behaviors through trial and error, adapting to dynamic environments and improving their performance over time. This method of learning is appealing because it allows systems to develop behaviors that are not explicitly programmed but rather learned through interaction with their environment, mirroring the way humans learn from experience and adjust their actions to achieve desired outcomes [33].

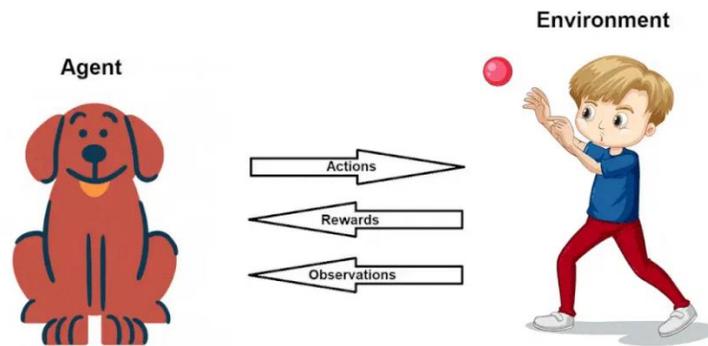


Figure 17: A visual representation of reinforcement learning [39].

3.6 Development of ANNs for Supervised Learning tasks

This study uses supervised learning for classification, and the following sections detail common techniques for training and evaluating machine learning models with labeled data.

3.6.1 Data Preprocessing

Collecting and preprocessing training data is crucial in machine learning applications. The amount and quality of available data heavily influences the final results, so considerable effort is dedicated to extracting and refining them. Preprocessing typically involves tasks like cleaning data (dealing with invalid entries), normalization, reducing features (through extraction or selection), and creating new features based on existing ones. This stage is essential for optimizing the data before feeding it into machine learning models to ensure accurate and efficient learning and prediction.

3.6.1.1 Missing Data Handling

Handling missing data is a critical component of data preprocessing. Missing data can arise from a variety of sources, such as errors in data entry, equipment malfunctions, or non-responses in surveys. The presence of missing data can significantly distort the results of analyses, leading to biased estimates and incorrect conclusions. Therefore, it is essential to employ effective strategies for handling and imputing missing data to ensure the integrity of the dataset and the validity of the analyses performed.

Deletion methods are one of the simplest ways to handle missing data. Listwise deletion removes entire records that contain any missing values, which can be effective when the amount of missing data is small. However, this method can lead to significant information loss and potential bias if the missing data is not randomly distributed. Pairwise deletion, on the other hand, only

uses available data for each specific analysis, thereby retaining more data but complicating the analysis process.

Imputation methods involve filling in missing values with substituted ones. Simple imputation techniques include replacing missing values with the mean, median, or mode of the observed data. While these methods are easy to implement, they can reduce data variability and introduce bias, particularly if the data distribution is skewed. More sophisticated imputation techniques, such as hot deck imputation, replace missing values with values from similar records within the dataset, maintaining data variability and relationships [40].

Regression imputation utilizes regression models to predict and fill in missing values based on other observed data. This method leverages the relationships between variables to provide more accurate imputations but assumes that these relationships are correctly specified. Another effective method is the K-Nearest Neighbors (KNN) imputation, which identifies 'k' similar instances (neighbors) and uses their values to impute the missing ones. This technique can handle complex relationships within the data but can be computationally intensive for large datasets. Additionally, machine learning models such as decision trees or random forests can predict missing values by learning from the observed data, capturing complex patterns but requiring significant computational resources and expertise [41].

3.6.1.2 Data Normalization

Normalization is widely adopted in applications involving ANNs due to its proven ability to improve performance and speed up training. By constraining input data to a predefined range, normalization reduces the impact of random outliers—values that don't represent typical data points in the training set and might result from measurement errors. This approach ensures that outlier values have minimal influence on the final outcomes. Moreover, normalization is crucial when input features vary widely in scale. By transforming these features, normalization ensures that all features contribute equally during the weight adjustment process from the start of training, preventing larger numerical values from overshadowing smaller ones.

Normalization techniques vary in their approach, but two widely used methods are Z-score normalization and Min-Max normalization. These methods are crucial for preparing data in machine learning, ensuring consistency and reducing the impact of varying scales across different features, thus optimizing model training and performance.

Z-score normalization adjusts data based on the mean and standard deviation of each feature, ensuring that values are standardized relative to the feature's distribution:

$$x'_{i,n} = \frac{x_{i,n} - \mu_i}{\sigma_i} \quad (25)$$

where σ_i and μ_i represent the average and standard deviation for the i input characteristic, respectively. Additionally, $x_{i,n}$ and $x'_{i,n}$ denote the initial and the final values of the n entry of the i feature.

On the other hand, Min-Max normalization scales data to a specified range, usually $[0, +1]$ or $[-1, +1]$, by mapping the original values between the minimum and maximum observed within each feature:

$$x'_{i,n} = \frac{x_{i,n} - \min(x_i)}{\max(x_i) - \min(x_i)} \cdot (nMax - nMin) + nMin \quad (26)$$

here, $\min(x_i)$ and $\max(x_i)$ indicate the smallest and largest values of the i feature. $nMin$ and $nMax$ refer to the lower and upper limits of the range that all the data will fall within after the normalization [42].

3.6.2 Training of the Model

The simplest training method involves splitting the available dataset into two distinct parts: the training set and the test set. This approach is known as the hold-out method. Only the training set is used to adjust the model's weights, while the test set (typically 20%-30% of the total dataset) is reserved for model evaluation based on selected metrics after training. This technique is popular due to its computational ease and relatively reliable results, especially with large datasets. However, with smaller datasets, it's often preferable to use all the data for training to ensure the model accounts for all possible scenarios. A downside of the hold-out method is its heavy dependence on how the data is split, which can impact reliability.

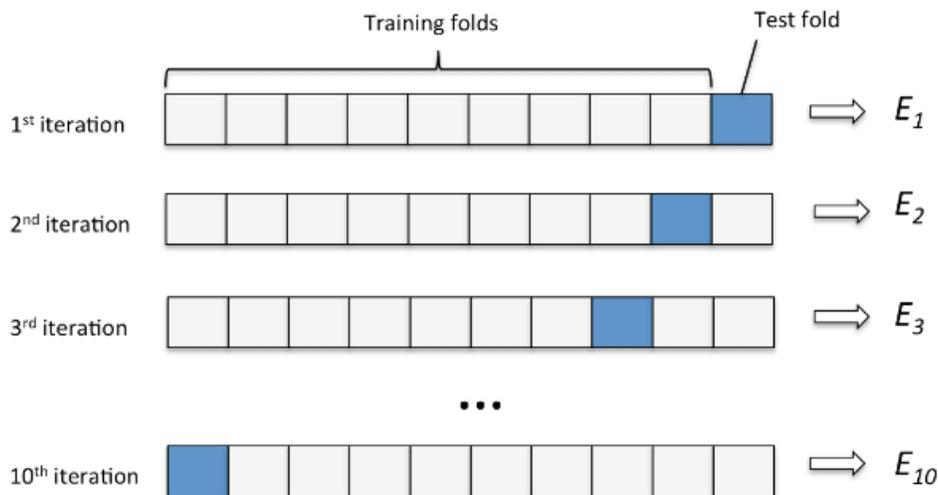


Figure 18: k-fold cross-validation for $k = 10$ [43].

To overcome these limitations, a widely-used alternative is cross-validation, specifically k-fold cross-validation (k-fold CV). This technique divides the data into k equal (or nearly equal) parts, or folds, and applies the hold-out method k times. Each time, a different fold is used as the test set, while the remaining k-1 folds form the training set. The model's error is averaged over the k iterations. This method trains and evaluates the model on k different sets, reducing the chance of performance variations due to data splitting. Additionally, it utilizes the entire dataset for weight adjustment, which is crucial for small datasets. However, k-fold CV requires more computational effort due to the increased number of repetitions [41].

Leave-one-out cross-validation is an extreme form of k-fold cross-validation where k equals the total number of samples (N). In leave-one-out cross-validation, the model is trained N times, each time with a training set of size N – 1 and evaluated on a single sample. The main advantage of this type of CV is its nearly complete use of the original dataset each time, though it requires significant time and computational resources. Another important variant is stratified k-fold cross-validation, which ensures that each subset accurately represents the original dataset by reorganizing data to match class distributions before splitting into k folds. This method is especially beneficial for datasets with imbalanced class distributions [44].

3.6.3 Evaluation of the Model

Once an ANN has completed training, evaluating its performance using specific metrics becomes essential. The choice of metrics depends on whether the application is focused on classification or regression tasks. For the scenario involving binary classification, where one category is designated as positive and the other as negative, key terms such as true positive (TP), true negative (TN), false positive (FP), and false negative (FN) are utilized.

TP indicates instances where the model correctly predicts a positive outcome, aligning with the actual positive values in the dataset, thereby ensuring accurate classification. TN describes instances where the model correctly predicts a negative outcome, aligning with the actual negative values, thus also ensuring accurate classification. FP occurs when the model incorrectly predicts a positive outcome for samples that actually belong to the negative class, leading to misclassification. Conversely, FN refers to instances where the model incorrectly predicts a negative outcome for samples that actually belong to the positive class, also resulting in misclassification. These terms underpin the following metrics tailored for binary classification:

- Accuracy measures the overall correctness of predictions by the model, calculated as the ratio of correct predictions (TP + TN) to the total number of samples:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (27)$$

- Precision quantifies the precision of the model's positive predictions, calculated as the ratio of true positive predictions (TP) to all positive predictions made by the model:

$$Precision = \frac{TP}{TP + FP} \quad (28)$$

- Recall (or sensitivity) gauges the model's ability to correctly identify all actual positive samples, calculated as the ratio of true positive predictions (TP) to all actual positive samples:

$$Recall = \frac{TP}{TP + FN} \quad (29)$$

- Specificity measures the model's ability to correctly identify all actual negative samples, calculated as the ratio of true negative predictions (TN) to all actual negative samples:

$$Specificity = \frac{TN}{TN + FP} \quad (30)$$

- F1-score balances precision and recall into a single metric, calculated as the harmonic mean of precision and recall:

$$F1_score = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} \quad (31)$$

- Area Under the Curve (AUC) from the ROC (Receiver Operating Characteristic) curve evaluates the classifier's ability to distinguish between classes. The ROC curve plots the true positive rate (Recall) against the false positive rate (FPR) for various classification thresholds. A higher AUC value signifies better classification performance, with 1 indicating perfect prediction and values above 0.5 showing that the model achieves more TP and TN than FP and FN [45].

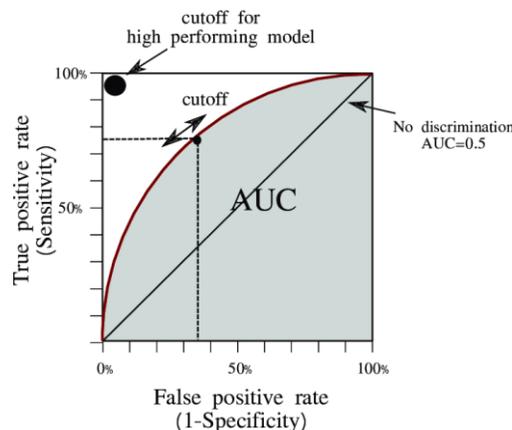


Figure 19: ROC curve and AUC [46].

- Youden's J statistic aids in finding the optimal classification threshold by maximizing the difference between true positive and false positive rates across thresholds:

$$J = \text{Sensitivity} + \text{Specificity} - 1 \quad (32)$$

This statistic provides a unified metric for assessing overall classifier performance. Maximizing Youden's J helps determine the optimal operating point on the ROC curve, ensuring the classifier achieves a balanced trade-off between sensitivity and specificity for the given task [47].

4. Related Work

A study of the existing related literature was conducted to achieve a deeper understanding and familiarity with the research problem being addressed. The rapid development of computational methods and the greater availability of relevant data, combined with the rapid spread of diabetes and the complexity of the disease, have attracted the interest of many researchers to this field. Consequently, the relevant literature has become rich and multidimensional. In this review, recent studies that utilize machine learning techniques, rather than other classical statistical methods, were concentrated on. As the focus was on assessing the risk of nocturnal hypoglycemia episodes, emphasis was placed on machine learning models that predict blood glucose levels and related events (hypoglycemia/hyperglycemia). Additionally, special attention was given to physiological system simulation models that can be integrated into prediction systems.

4.1 Physiological Prediction Models

Over the past few decades, numerous studies have introduced models of insulin action and glucose kinetics, utilizing experimental data to measure glucose production, glucose utilization, and the absorption of insulin and meals. Predominantly, these models are compartmental models used for the dynamic analysis of substance kinetics in physiological systems. Among the most renowned models for simulating insulin and glucose kinetics are the Dalla Man model, the Hovorka model, and the Bergman model. These systems enable the computation of various parameters, including subcutaneous insulin absorption, gastric emptying rate, carbohydrate digestion and absorption, insulin kinetics, and glucose metabolism [48].

Specifically, the Dalla Man model comprises a glucose subsystem and an insulin subsystem that are interconnected by the regulation of insulin on glucose utilization and endogenous production. On the other hand, the Bergman model utilizes a three-compartment model to represent the concentrations of plasma insulin, exogenous insulin, and plasma glucose. The Hovorka model, in contrast, employs a two-compartment model to describe the absorption of subcutaneous insulin and the absorption of glucose from the intestine [49-51].

Additionally, to evaluate glucose absorption at the cellular level, the activation process of the intracellular insulin signaling pathway can be simulated. A widely recognized simulation system for this purpose is the model by Sedaghat. This model, based on insulin concentration and utilizing a system of 20 differential equations, estimates the proportion of GLUT4 transporters on the cell membrane. Additionally, this model has been expanded and enhanced through more recent related research [52, 53].

4.2 Machine Learning Models for Predicting Blood Glucose Levels

The 2015 study by Zarkogianni et al. focuses on evaluating machine learning models for predicting future glucose levels in individuals with T1D. The research utilized data from 10 real T1D patients (7 males and 3 females), incorporating CGM data and physical activity data as inputs. Two sets of input data were used: the first included the most recent glucose measurement and CGM-derived glucose change, while the second integrated these with physical activity data. Four machine learning models were tested: Feedforward Neural Network (FNN), Self-Organizing Map (SOM), Wavelet-Transformed Feedforward Neural Network (WFNN), and Linear Regression Model (LRM). These models aimed to predict future glucose values at 30, 60, and 120-minute horizons with updates every 5 minutes, using a personalized training approach for each patient.

Evaluation employed a 10-fold cross-validation method, assessing model performance through metrics such as root-mean-square error (RMSE), correlation coefficient (CC), mean absolute relative difference (MARD), and continuous glucose-error grid analysis (CG-EGA). The study highlighted that while all models performed well within the euglycemic range, the SOM-based model excelled in predicting hypoglycemic and hyperglycemic events. It achieved superior RMSE, CC, and MARD values for both sets of input data, underscoring its efficacy in glucose prediction for individuals with T1D [54].

Published in 2020, the study by Cappon et al. utilized data from 6 real patients sourced from the second version of the OhioT1DM dataset, focusing on integrating various input features including CGM measurements, injected insulin, reported meals, physical activity, and correction boluses. The research employed a bidirectional LSTM neural network architecture designed to predict future blood glucose levels at 30-minute and 60-minute horizons for each patient, reflecting a personalized approach where individual models were trained for different subjects.

Model performance was evaluated using metrics such as RMSE, mean absolute error (MAE), and time gained (TG). The study emphasized the interpretability of its findings, noting that high CGM values predicted correspondingly high blood glucose levels, while elevated insulin levels negatively influenced blood glucose dynamics, mirroring actual clinical scenarios. Additionally, the model effectively captured the impact of meal intake on blood glucose levels, demonstrating its capability to adjust predictions based on real-time physiological responses. Despite its achievements, the study acknowledged limitations, including the absence of comparator models for benchmarking bidirectional LSTM performance and the relatively small dataset size [55].

In 2022, Mosquera-Lopez et al. conducted a study using data from 250 real patients with Type 1 Diabetes sourced from the Tidepool Big Data Donation Dataset. Their research focused on enhancing predictive modeling by integrating comprehensive measures of glucose variability. They employed a hybrid approach combining an equation for estimating insulin on board (IOB) with LSTM neural networks. This model was specifically tailored to incorporate CGM readings and scaled IOB data from the preceding 3 hours, addressing crucial mid- and short-term

dependencies affecting glucose levels. The study rigorously compared the performance of the hybrid LSTM model against various benchmarks, including naive approaches, Ridge linear regression, and Random Forest models. Notably, the LSTM model exhibited superior predictive accuracy, particularly in scenarios tailored to different insulin therapies reflecting real-world treatment variability among Type 1 Diabetes patients.

Evaluation utilized a hold-out testing dataset and encompassed a comprehensive set of performance metrics, including traditional measures like RMSE and MAE, as well as novel indices like the Glucose Variability Impact Index (GVII) and Glucose Prediction Consistency Index (GPCI). Key findings underscored the LSTM model's robust performance in accurately forecasting glucose values across diverse insulin therapy contexts. However, the study also identified challenges, such as the model's tendency to overestimate hypoglycemia and underestimate hyperglycemia events, highlighting ongoing complexities in accurately predicting extreme glucose fluctuations. Moreover, the research emphasized the significant influence of glucose data variability on prediction outcomes, emphasizing the critical need for model refinements to effectively manage these variations and improve overall predictive accuracy in Type 1 Diabetes management [56].

Published in 2023, the study by Toffanin et al. delves into the development of an Enhanced Personalized LSTM (EP-LSTM) model aimed at predicting blood glucose levels and preventing hypoglycemia and hyperglycemia in Type 1 Diabetes patients. Using data from 100 in silico adult patients from the UVA/Padova dataset, the research incorporates inputs such as insulin injected through a pump, patient-reported meal intake, and past CGM values from 40 minutes prior to the prediction horizon.

The EP-LSTM model's performance was assessed using a range of metrics, including RMSE, index of fitting (FIT), downward delay (DD), upward delay (UD), true positive rate, positive predictive value (PPV), and F1 score. The study reported an RMSE of 6.45 and a FIT of 79.40%, with an average delay of approximately 9 minutes. For hypoglycemia detection, the alarm system achieved a TPR of 76.92%, PPV of 83.33%, and an F1 score of 78.79%. In hyperglycemia detection, it reached a TPR of 89.13%, PPV of 85.29%, and an F1 score of 83.87%. These results indicate that the EP-LSTM model performs effectively in predicting and preventing extreme glucose events, underscoring its potential utility in personalized diabetes management [57].

4.3 Machine Learning Models for Predicting Nocturnal Hypoglycemia

In the aforementioned studies, glucose prediction is treated as a regression problem, where the objective is to predict continuous numerical values. By focusing specifically on hypoglycemia, the problem can be transformed into a classification problem aimed at predicting hypoglycemic episodes. The following examines some research that explores this specific approach.

Published in 2020, the study by Jensen et al. investigates the use of a linear discriminant analysis (LDA) classifier to predict nocturnal hypoglycemia in individuals with Type 1 Diabetes. The

research utilizes data from 463 real patients from the Onset 5 trial by Novo Nordisk A/S, integrating CGM data, meal intake, and bolus insulin data after feature extraction. The LDA classifier is personalized, trained on the previous three days of patient recordings to predict nocturnal hypoglycemia effectively.

The model's performance is evaluated using metrics such as ROC-AUC, specificity, and sensitivity, achieving a ROC-AUC of 0.79, a sensitivity of 75%, and a specificity of 70%. The average prediction horizon is approximately three hours and 15 minutes before a hypoglycemic event, typically occurring around 3:15 am, with warnings issued at midnight. Despite the promising results, the study acknowledges significant limitations, including the absence of basal insulin data and physical activity measures, which are known contributors to nocturnal hypoglycemia. Furthermore, the paper lacks detailed technical information, highlighting areas for future research improvement [58].

In 2020, Mosquera-Lopez et. al conducted a study that explores the use of support vector regression (SVR) models to predict and prevent nocturnal hypoglycemia in individuals with Type 1 Diabetes. The research utilizes data from both real and in silico patients, including 124 real insulin pump users from the Tidepool Big Data Donation Dataset, along with additional datasets for algorithm validation. Thirteen features were extracted from CGM, insulin, and meal data. The SVR model is optimized using decision theory to maximize the benefit of accurate hypoglycemia prediction while minimizing the cost of inaccurate predictions. This decision support tool helps patients make optimal decisions about whether to consume carbohydrates at bedtime to prevent nocturnal hypoglycemia. The SVR model was compared against a simple bedtime glucose heuristic, which advises patients to consume carbohydrates if their bedtime glucose is below 8.28 mmol/L.

The model's performance was evaluated with metrics including AUC, sensitivity, specificity, Pearson correlation, and RMSE between actual and predicted minimum glucose levels during the night. The SVR model outperformed the simple bedtime glucose heuristic in terms of specificity for the same sensitivity value. Recommendations based on SVR predictions could potentially result in 2.5 overtreated cases per month but applying the SVR algorithm earlier or later in the night yielded lower accuracy. The study found that a variable carbohydrate intervention based on SVR predictions could reduce nocturnal hypoglycemia by up to 77% without affecting overall time in the target glucose range. This indicates that the SVR algorithm enhances the specificity of predictions and increases time in the target glucose range, demonstrating its potential to effectively manage nocturnal hypoglycemia in Type 1 Diabetes patients [59].

Published in 2022, the study by Parcerisas et al. leverages data from both real and in silico patients, including 10 real patients from a clinical trial dataset (Hospital Clinic de Barcelona NCT03711656) and in silico patients from a modified version of the UVA Padova simulator. The

research extracts 17 features from CGM readings and activity tracker data, employing bolus on board, carbohydrate on board, and activity on board models for feature extraction. Various machine learning algorithms were tested, including ANNs, multinomial naive Bayes, AdaBoost, support vector machines (SVM), LDA, and LSTM networks, with SVM demonstrating the best performance. The model's objective is to predict the risk of nocturnal hypoglycemic events, utilizing both population-based and personalized approaches.

Performance evaluation involved metrics such as AUC-ROC, sensitivity, specificity, Matthews' correlation coefficient (MCC), F1 score, and geometric mean (Gmean). The results showed minimal differences between population and personalized models in terms of median sensitivity and specificity. However, the population model exhibited superior F1 scores and Gmean metrics. The introduction of rescue carbohydrates significantly reduced the number of nocturnal hypoglycemic events, with 30 grams identified as the optimal amount for prevention. The study concludes that the population model's performance is on par with individual models, underscoring its potential effectiveness in minimizing nocturnal hypoglycemia in Type 1 Diabetes patients [60].

In 2022, Berikov et al. conducted a study using a dataset of 406 real adult patients undergoing basal bolus insulin therapy, which presented an imbalance in data distribution. The study's input variables included CGM readings and 23 clinical and laboratory parameters such as age, sex, BMI, diabetes duration, insulin regimen details, and various health conditions. The models employed encompassed Random Forest with 500 trees, Logistic Linear Regression with Lasso regularization (LogRLasso), and ANNs optimized using the Levenberg–Marquardt algorithm. These models were developed with the goal of predicting hypoglycemia risk through a personalized approach.

Evaluation of model performance focused on metrics like ROC-AUC, sensitivity, and specificity. Notably, LogRLasso models using only CGM data achieved the highest AUC values. Random Forest demonstrated superior accuracy when integrating CGM with clinical data, while ANNs slightly underperformed when trained on CGM alone or in combination with clinical variables. Clinical insights from the study highlighted positive correlations between insulin dose, diabetes duration, and proteinuria with hypoglycemia risk, while HbA1c, estimated glomerular filtration rate (eGFR), and BMI showed negative associations. These findings underscore the potential of machine learning to enhance the prediction and management of nocturnal hypoglycemia among hospitalized patients with Type 1 Diabetes [61].

5. Development of attention-based LSTM models for the prediction of nocturnal hypoglycemia events in T1DM

The preceding sections aimed to provide the foundational theoretical background into the tools applied in this study's practical implementation. Following this introduction is the experimental phase, which centers on a classification task using glucose levels and other variables to forecast the likelihood of imminent nocturnal hypoglycemic events. This chapter includes a description of available data, and the methodologies used for their processing, and an outline of the model and methodologies utilized for predictions and result interpretation. Notably, Python version 3.10.12 was the implementation platform, leveraging essential libraries such as Keras, scikit-learn, NumPy, pandas, Matplotlib and GEKKO for assembling training datasets and refining the model through development and evaluation phases.

5.1 Dataset and Data Preparation

5.1.1 OhioT1DM Dataset

The model was developed and evaluated using data derived from actual patients diagnosed with Type 1 Diabetes, sourced from the OhioT1DM Dataset. This dataset serves as a pivotal resource aimed at fostering advancements in the prediction of blood glucose levels. Over an eight-week period, it captured a comprehensive array of continuous CGM data, insulin administration records, physiological sensor readings, and self-reported life events for 12 individuals diagnosed with type 1 diabetes.

Initially launched in 2018 for the first Blood Glucose Level Prediction Challenge with data from six participants, the dataset has since expanded to include an additional six individuals for the 2020 challenge. Each participant in the OhioT1DM Dataset was assigned a randomly generated ID to safeguard anonymity. Throughout the data collection period, all contributors utilized insulin pump therapy alongside CGM devices from Medtronic (models 530G or 630G). They also reported life-event data via a bespoke smartphone application and provided physiological metrics from either Basis Peak or Empatica Embrace fitness bands. Key dataset features included CGM readings every five minutes, periodic self-monitored blood glucose values, detailed logs of insulin doses (both basal and bolus), self-reported mealtimes with carbohydrate estimates, and timestamps for activities such as exercise, sleep, work, stress, and illness. Participants that were using the Basis Peak band contributed additional data points every five minutes on heart rate, galvanic skin response (GSR), skin temperature, air temperature, and step count, while those that were using the Empatica Embrace band provided GSR, skin temperature, and minute-by-minute acceleration data. Sleep patterns and subjective sleep quality assessments were also logged where applicable.

ID	BGLP Challenge	Gender	Age	Training	Testing
540	2020	Male	20-40	11947	2884
544	2020	Male	40-60	10623	2704
552	2020	Male	20-40	9080	2352
567	2020	Female	20-40	10858	2377
584	2020	Male	40-60	12150	2653
596	2020	Male	60-80	10877	2731
559	2018	Female	40-60	10796	2514
563	2018	Male	40-60	12124	2570
570	2018	Male	40-60	10982	2745
575	2018	Female	40-60	11866	2590
588	2018	Female	40-60	12640	2791
591	2018	Female	40-60	10847	2760

Table 1: Cohort, gender, age group, training data and testing data entries for each contributor of the OhioT1DM dataset [62].

To facilitate comprehensive research and analysis, the dataset was meticulously organized into XML files for each participant, distinguishing between training and testing datasets. This structured approach allows researchers to develop and refine predictive models for blood glucose levels, ensuring robustness and applicability in clinical settings [62].

Here, to develop this model, CGM, basal insulin, temporary basal insulin, bolus insulin, and meal data for each patient were utilized. As mentioned above (sections 2.1 and 2.5), blood glucose levels and the overall progression of diabetes are greatly influenced by diet and insulin administration.

5.1.2 Data Preparation

The objective of this study was to predict nocturnal hypoglycemia events. However, numerous gaps in self-reported sleep event data posed a significant challenge. To address this issue, a fixed sleep window from 20:00 to 08:00 the next morning was hypothesized, establishing a 12-hour prediction horizon.

The aim was then to identify nocturnal hypoglycemia events for each patient. A hypoglycemic event was defined as a CGM measurement below 70 mg/dL. If such an event occurred within the designated sleep window, the night was labeled as a nocturnal hypoglycemia night (label = 1); otherwise, it was labeled as a non-hypoglycemia night (label = 0). The data analysis yielded the following results:

ID	BGLP	Data Split	Data Entries	CGM Missing Values	Basal Values	Temp_basal Values	Bolus Values	Meal Values	Noc Hypos
559	2018	train	11808	1300	163	34	152	150	17
		test	2592	362	32	6	36	29	3
563	2018	train	12960	1008	87	2	347	129	12
		test	2304	121	12	0	89	27	0
570	2018	train	11520	826	118	3	326	136	8
		test	2592	135	29	0	84	33	1
575	2018	train	12960	1238	126	12	187	243	21
		test	2304	128	22	1	36	45	4
588	2018	train	12960	608	74	34	182	221	11
		test	2592	89	8	3	40	37	1
591	2018	train	12672	2113	113	33	261	212	21
		test	2592	87	29	5	51	41	4
540	2020	train	12960	1302	24	232	309	73	26
		test	2592	170	8	55	87	27	5
544	2020	train	12384	2049	103	43	134	159	11
		test	2880	420	32	12	39	38	0
552	2020	train	10944	2152	45	16	336	78	9
		test	3456	1586	18	7	102	21	2
567	2020	train	13248	2678	135	45	313	32	14
		test	2592	482	14	9	54	0	3
584	2020	train	12960	1098	27	12	268	95	7
		test	2592	330	20	0	54	23	2
596	2020	train	14112	3412	23	7	208	265	15
		test	2592	260	31	1	38	54	1

Table 2: Mapping of the OhioT1DM dataset.

Initially, the development of personalized models for each of the 12 patients was considered. However, the extensive input space combined with the limited training data resulted in models with poor learning performance, indicating an inability to effectively train on the available data. Consequently, a population-based model was developed instead.

Given the prevalence of missing CGM values and the potential bias introduced by interpolating these gaps, days with more than 24 consecutive missing CGM values were excluded from the analysis. The initial training and testing datasets comprised 526 and 110 days of data, respectively, with nocturnal hypoglycemia rates of 0.327 and 0.236, respectively. After exclusions, the final training and testing datasets were reduced to 352 and 77 days, with nocturnal hypoglycemia rates of 0.321 and 0.195, respectively.

Data Split	Data Entries	Data Days	Noc Hypos	Noc Hypos Rate
Initial Train	151,488	526	172	0.327
Final Train	101,376	352	113	0.321
Initial Test	31,680	110	26	0.236
Final Test	22,176	77	15	0.195

Table 3: Global datasets with the combined data for each patient.

To create time series data for the model, 24-hour sequences of CGM readings, carbohydrate meals, and insulin doses were compiled, starting from 20:00 (t_{start}) on the first day and ending at 20:00 (t_{end}) the following day, resulting in time series of $[t_{start}, t_{end}]$. Since CGM recordings occurred every 5 minutes, each 24-hour sequence consisted of 288 values (24 hours x 12 measurements per hour). This data was then used to predict the probability of nocturnal hypoglycemia within the sleep window from 20:00 (t_{end}) to 08:00 ($t_{end} + 12h$) the next morning.

5.2 Architecture of Hybrid Predictive Model

To predict the risk of nocturnal hypoglycemic events, a hybrid approach combining physiological system simulation models and machine learning techniques was employed. This method aimed to accurately model the (i) absorption of glucose from meals into the bloodstream, (ii) insulin absorption from subcutaneous tissue to bloodstream, and (iii) insulin signaling activation within cells. The outputs from these simulation models, along with CGM time series data, were integrated using an LSTM-based neural network for predictive analysis.

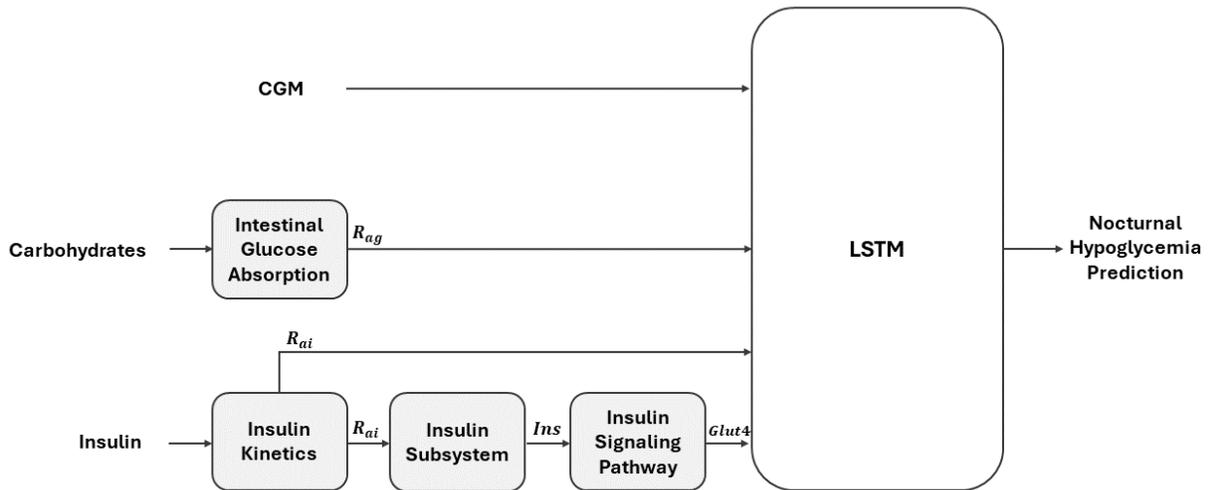


Figure 20: Diagrammatic representation of the hybrid model for predicting nocturnal hypoglycemic events.

Central to the methodology were compartmental models, which are crucial for studying substance kinetics within physiological systems. These models accommodate both exogenous substances (such as drugs) and endogenous ones (such as hormones), detailing processes like

production, distribution, transport, and utilization of these substances. Each compartment within these models represents a quantity of substance that behaves as if it is uniformly mixed and kinetically homogeneous, with interconnected pathways defining substance flow between compartments [63]

5.2.1 Intestinal Glucose Absorption Compartment Model

For more precise utilization of available carbohydrate information, the Dalla Man compartmental model was used. This model illustrates the glucose flow in the stomach and intestine and is composed of three compartments: two for the stomach and one for the intestine [50]. The associated system of differential equations is provided below:

$$Q_{sto}(t) = Q_{sto1}(t) + Q_{sto2}(t) \quad (33)$$

$$Q_{sto}(0) = 0 \quad (34)$$

$$\dot{Q}_{sto1}(t) = -k_{gri} \cdot Q_{sto1}(t) + D \cdot \delta(t) \quad (35)$$

$$Q_{sto1}(0) = 0 \quad (36)$$

$$\dot{Q}_{sto2}(t) = -k_{empt} \cdot Q_{sto2}(t) + k_{gri} \cdot Q_{sto1}(t) \quad (37)$$

$$Q_{sto2}(0) = 0 \quad (38)$$

$$\dot{Q}_{gut}(t) = -k_{abs} \cdot Q_{gut}(t) + k_{empt} \cdot Q_{sto2}(t) \quad (39)$$

$$Q_{gut}(0) = 0 \quad (40)$$

$$R_{ag}(t) = \frac{f \cdot k_{abs} \cdot Q_{gut}(t)}{BW} \quad (41)$$

$$R_{ag}(0) = 0 \quad (42)$$

$$k_{empt}(Q_{sto}) = k_{\min} + \frac{k_{\max} - k_{\min}}{2} \cdot (\tanh(a \cdot (Q_{sto} - b \cdot D)) + 1) \quad (43)$$

$$a = \frac{5}{2 \cdot D \cdot (1 - b)} \quad (44)$$

where:

- Q_{sto} (mg) was the total amount of glucose in the stomach,
- Q_{sto1} (mg) was the amount of glucose in solid form in the stomach,
- k_{gri} (min^{-1}) was the grinding rate of glucose,
- D (mg) was the amount of carbohydrates ingested,
- $\delta(t)$ was the Dirac delta function,
- Q_{sto2} (mg) was the amount of glucose in liquid form in the stomach,
- k_{empt} (min^{-1}) was the rate of gastric emptying,
- Q_{gut} (mg) was the amount of glucose in the intestine
- k_{abs} (min^{-1}) was the absorption rate of glucose from the intestine
- R_{ag} (mg/kg/min) was the rate of appearance of glucose in the blood,
- f was the fraction of the intestinal absorption which actually appeared in the plasma, and
- BW (kg) was the weight of the individual.

The constants k_{gri} , k_{abs} , f , b , k_{min} , and k_{max} were defined values found in the literature [64]. Thus, with a given quantity of carbohydrates D , the rate at which glucose appears in the blood was estimated. This information is clinically significant and can help improve the accuracy of predictions.

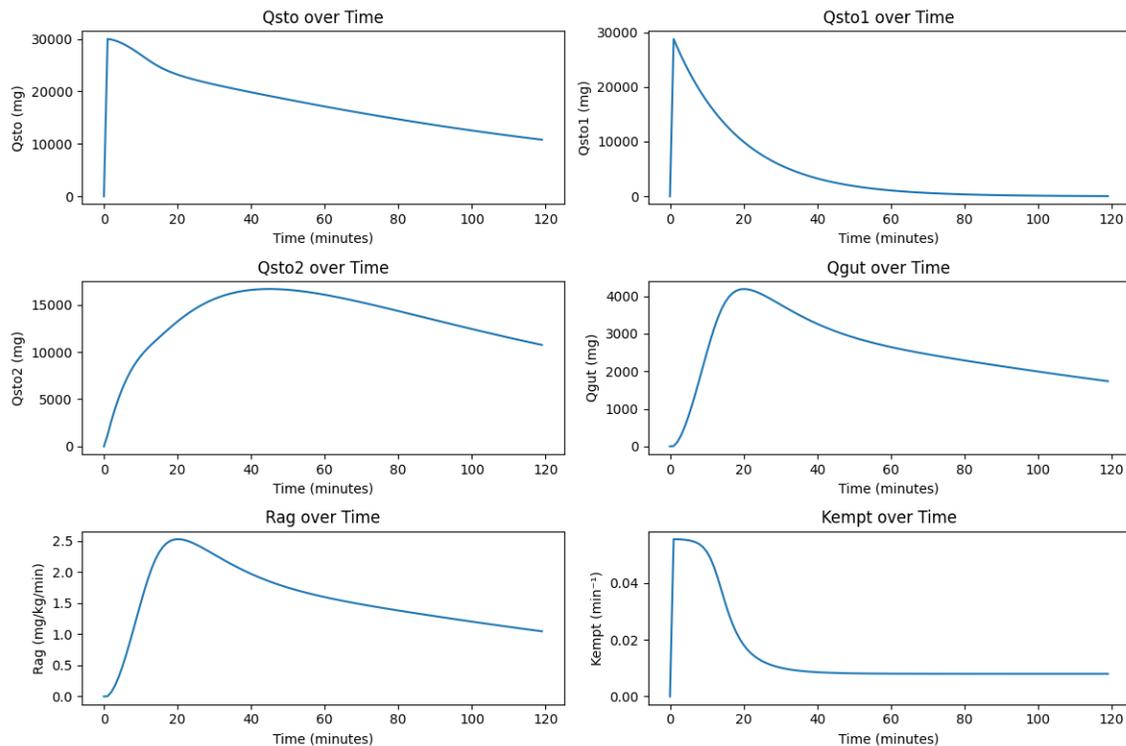


Figure 21: Computation of the values for Q_{sto} , Q_{sto1} , Q_{sto2} , Q_{gut} , R_{ag} and k_{empt} over a 120-minutes period, given a carbohydrate intake of 30 grams and a body weight of 85 kilograms, utilizing the Intestinal Glucose Absorption CM.

It was observed that the curve of the rate of glucose appearance in the blood, $R_{ag}(t)$, was shaped by the total intake of meals consumed recently, depending on the amount of carbohydrates ingested and the timing of each meal. Hence, it was concluded that $R_{ag}(t)$ represented a more substantial and accurate depiction of the overall glucose absorption process, and thus, it was used as an input feature for improved predictions.

Here, given a glucose time series that outlined a patient's glycemic profile within the 24-hour period $[t_{start}, t_{end}]$, the recorded carbohydrate intakes within this time interval were considered. Using this data, $R_{ag}(t)$ was calculated for $t \in [t_{start}, t_{end} + 12h]$ by applying the Intestinal Glucose Absorption CM and superimposing the curves generated by the model for each specific carbohydrate intake.

The simulation was conducted over a 36-hour period: 24 hours corresponding to the recorded carbohydrate intakes, followed by an additional 12 hours to predict nocturnal glucose absorption rate. Although $R_{ag}(t)$ was computed for the entire simulation period, only the last 12 hours were retained as input features for the prediction of the nocturnal hypoglycemia probability. The $R_{ag}(t)$ values were aggregated every 5 minutes to maintain synchronization with the CGM data. Therefore, the input features for the LSTM-based model included the $R_{ag}(t)$ time series vector, which contained 144 values (12 hours x 5 measurement per hour).

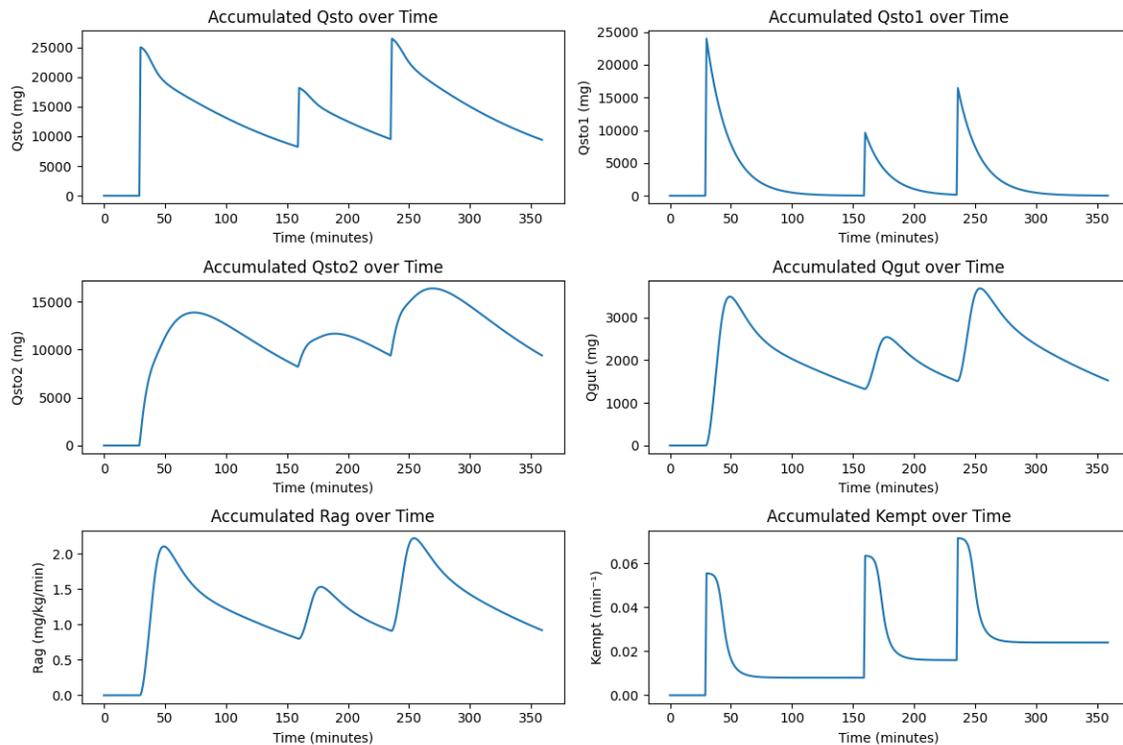


Figure 22: Computation of the values for Q_{sto} , Q_{sto1} , Q_{sto2} , Q_{gut} , R_{ag} and k_{empt} for three different carbohydrate intakes of 25g, 10g and 17g over a 360-minutes period. The calculation was conducted using the Intestinal Glucose Absorption CM and by superimposing the corresponding output curves.

5.2.2 Blood Insulin Absorption Compartment Model

To represent the complex processes associated with the absorption of administered insulin from the subcutaneous space into the bloodstream, the following compartmental model of subcutaneous insulin kinetics was utilized [65]:

$$I_{sc1}'(t) = -(k_d + k_{a1}) \cdot I_{sc1}(t) + u(t) \quad (45)$$

$$I_{sc2}'(t) = k_d \cdot I_{sc1}(t) - k_{a2} \cdot I_{sc2}(t) \quad (46)$$

$$R_{ai}(t) = k_{a1} \cdot I_{sc1}(t) + k_{a2} \cdot I_{sc2}(t) \quad (47)$$

$$I_{sc1}(0) = I_{sc1ss} \quad (48)$$

$$I_{sc2}(0) = I_{sc2ss} \quad (49)$$

where:

- I_{sc1} ($pmol/kg$) represented the amount of non-monomeric insulin in the subcutaneous space,
- I_{sc2} ($pmol/kg$) denoted the amount of monomeric insulin in the same area,
- the rate of exogenous insulin infusion was indicated by $u(t)$ ($pmol/kg/min$),
- the constant rate of insulin degradation was given by k_d (min^{-1}), and
- k_{a1} (min^{-1}) and k_{a2} (min^{-1}) were the absorption rates for non-monomeric and monomeric insulin, respectively,
- R_{ai} ($pmol/kg/min$) described the rate of insulin appearance in the blood plasma.

The constants I_{sc1ss} and I_{sc2ss} were defined values found in the literature [65]. Overall, the described relationships outlined the diffusion process of administered insulin from the subcutaneous space into the blood plasma. It is important to note that insulin is typically administered in its more stable, polymeric form, which is then broken down into monomeric insulin that is absorbed more rapidly by the body.

Similarly to the case of Intestinal Glucose Absorption CM, given a glucose time series that outlined the glycemic profile over a 24-hour period $[t_{start}, t_{end}]$, the insulin doses administered during this interval, both bolus and basal, were considered. Additionally, the basal insulin doses recorded during the fixed 12-hour sleep window $[t_{end}, t_{end} + 12h]$ were included, as basal

insulin continued to be administered during sleep. Using this data, the rate of insulin appearance in the blood, $R_{ai}(t)$, was calculated over the entire period $t \in [t_{start}, t_{end} + 12h]$.

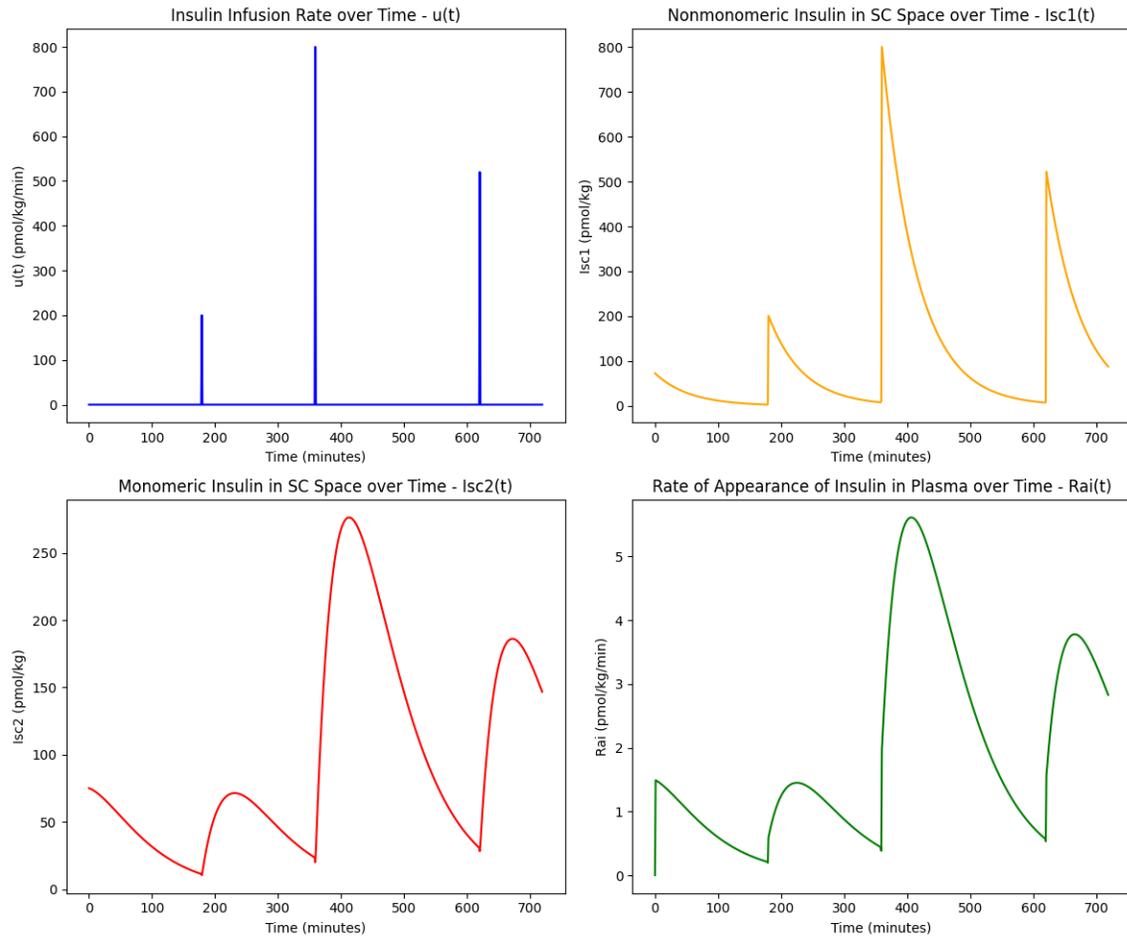


Figure 23: Computation of the quantities I_{sc1} , I_{sc2} and R_{ai} for three different insulin doses of 2.5, 10 and 6.5 units over a 720-minute period and a body weight of 75 kg, using the Blood Insulin Absorption CM.

The simulation for the $R_{ai}(t)$ was conducted over a 36-hour period, with only the last 12 hours retained as input features for prediction. To ensure synchronization with CGM data, the $R_{ai}(t)$ values were aggregated every 5 minutes. This results in a $R_{ai}(t)$ time series vector that contained 144 values, which was used as input in the LSTM-based model.

5.2.3 Insulin Kinetics Compartment Model

To examine the process of insulin absorption at the cellular level, two additional simulation models were used. The first model employed compartmental modeling, consisted of two compartments, liver and blood plasma, and aimed to calculate the insulin concentration in the plasma I , given the corresponding rate of appearance R_{ai} [50]:

$$I_P \dot{(t)} = -(m_2 + m_4) \cdot I_P(t) + m_1 \cdot I_L(t) + R_{ai}(t) \quad (50)$$

$$I_L \dot{(t)} = -(m_1 + m_3) \cdot I_L(t) + m_2 \cdot I_P(t) \quad (51)$$

$$I(t) = I_P(t)/V_I \quad (52)$$

$$I_L(0) = I_{Lb} \quad (53)$$

$$I_P(0) = I_{Pb} \quad (54)$$

where:

- I_P (pmol/kg) and I_L (pmol/kg) represented the amounts of insulin in the plasma and liver, respectively,
- I (pmol/L) denoted the insulin concentration in the plasma,
- R_{ai} (pmol/kg/min) was the rate of insulin appearance in the plasma,
- V_I (L/kg) denoted the volume of insulin distribution
- the constants m_1 (min^{-1}), m_2 (min^{-1}), m_3 (min^{-1}), and m_4 (min^{-1}) described the processes of insulin removal from the liver and peripheral tissues,
- I_{Lb} and I_{Pb} were the basal state values of I_L and I_P , respectively. These basal state values were found in the literature [50].

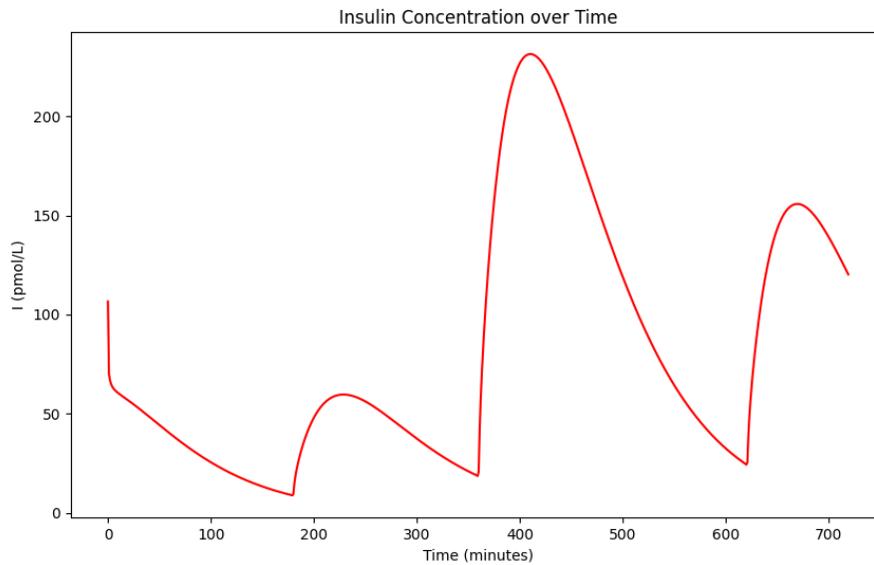


Figure 24: Computation of insulin plasma concentration $I(t)$ for three different insulin doses of 2.5, 10 and 6.5 units over a 720-minute period and a body weight of 75 kg, using the Insulin Kinetics CM and Blood Insulin Absorption CM (for $R_{ai}(t)$).

In this approach, the rate of insulin appearance R_{ai} was derived from the Blood Insulin Absorption CM, based on the rate of exogenous infusion of insulin.

5.2.4 Insulin Signaling Pathway Model

The mathematical model applied for simulating the activation of the intracellular insulin signaling pathway was the Sedaghat et al. model [52]. This model aimed to simulate key aspects of the PI3K/AKT insulin signaling pathway. The PI3K/AKT pathway is activated within a cell when insulin binds to its respective transmembrane receptors, initiating a cascade of protein phosphorylation. These intracellular protein interactions enable the signal to propagate from the receptor in various directions, ultimately causing the migration of vesicles containing the GLUT4 glucose transporters from the intracellular reservoir to the cell's plasma membrane. Any disruption in this signaling pathway can result in insulin resistance and decreased glucose uptake by cells.

The Sedaghat model used for this simulation involved 21 variables (x_1, x_2, \dots, x_{21}) derived from a system of 20 differential equations. The input x_1 represented the concentration of insulin in the interstitial fluid, while the variable x_{21} indicated the proportion of GLUT4 that reached the cell membrane. To estimate the concentration of insulin in the interstitial fluid, the plasma insulin concentration from the Insulin Kinetics CM was used, and it was assumed that the interstitial fluid insulin concentration was approximately 40% lower than in the plasma, as was reported in related studies [66].

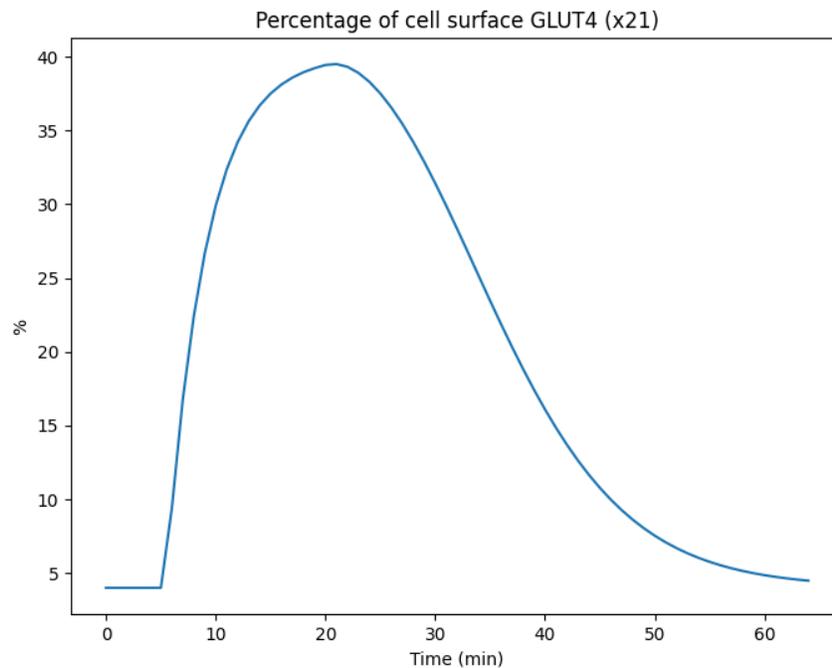


Figure 25: Estimation of the percentage of cell surface GLUT4 transporters using the Sedaghat simulation model, following the administration of insulin at a constant concentration of 10^{-5} M, over the time interval from $t = 0$ to $t = 15$ minutes.

In summary, to incorporate information related to the proportion of GLUT4 on the cell membrane (x_{21}) following insulin administrations, the Blood Insulin Absorption CM, the Insulin Kinetics CM and the Sedaghat model were utilized. Specifically, with a given glucose time series within the 24-hour interval $[t_{start}, t_{end}]$, the basal and bolus insulin doses administered by the patient during this period, along with only the basal doses for the period $[t_{end}, t_{end} + 12h]$ were considered.

The rate of insulin appearance R_{ai} was calculated for the period $[t_{start}, t_{end} + 12h]$ using the Blood Insulin Absorption CM. This rate was then used as input to estimate the insulin concentration I over the same period, employing the Insulin Kinetics CM. Finally, the percentage of GLUT4 receptors on the cell membrane (x_{21}) for the period $[t_{end}, t_{end} + 12h]$ was calculated using the Sedaghat model, with the insulin concentration I from this period as input. Thus, the estimated x_{21} was influenced by the insulin administrations during the preceding $[t_{end} - 24h, t_{end}]$ period. To ensure synchronization with CGM data, the x_{21} values were aggregated every 5 minutes. This results in a x_{21} time series vector that contained 144 values, which was used as input in the LSTM-based model.

5.2.5 Long Short Memory Neural Networks

To address the binary classification problem, a RNN was developed, consisting of four LSTM layers, a custom attention mechanism, a concatenation layer, a dropout layer, and a dense output layer. Each input vector— CGM , R_{ag} , R_{ai} , and $Glut_4$ —was processed by its corresponding LSTM layer. The outputs from these LSTM layers were then passed through custom attention layers. The attention-augmented outputs were concatenated, followed by dropout regularization, and finally fed into a dense layer for the final binary classification.

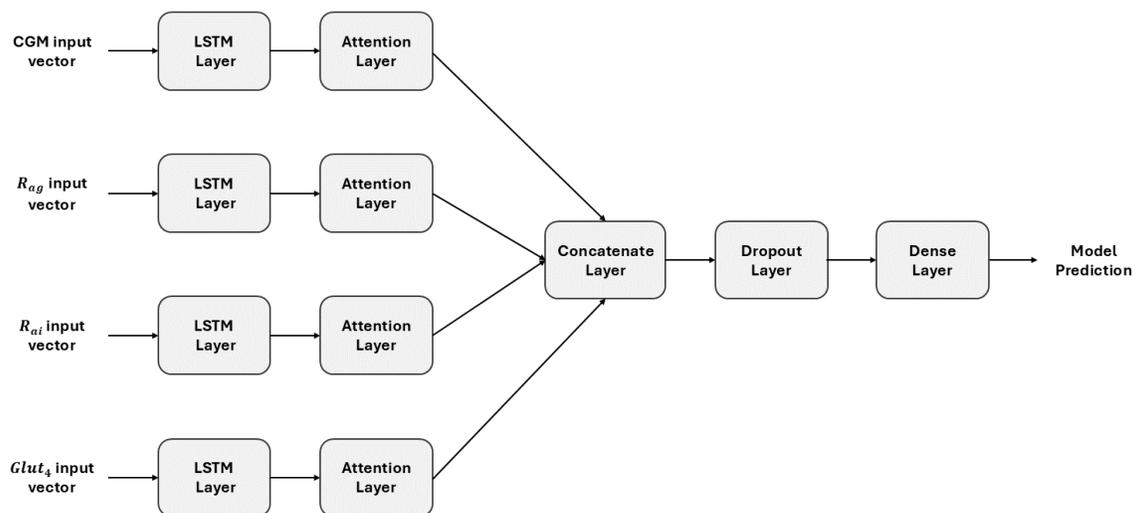


Figure 26: LSTM-based Model Architecture.

The custom attention mechanism was designed to enhance the performance of the LSTM layers by dynamically emphasizing the most relevant aspects of the input sequences. This mechanism began with a dense layer containing a single neuron and a tanh activation function, which computed raw attention scores for each timestep of the LSTM output. These scores were then flattened and normalized using a softmax activation function, resulting in attention weights. These weights were reshaped and applied via element-wise multiplication to the LSTM outputs, producing weighted outputs. Finally, a summation operation aggregated these weighted outputs across timesteps, creating a context vector that encapsulated the most relevant information from the input sequence. This attention mechanism, inspired by the work of Raffel et al. [67], allowed the network to focus on crucial aspects of the input, thereby improving predictive performance.

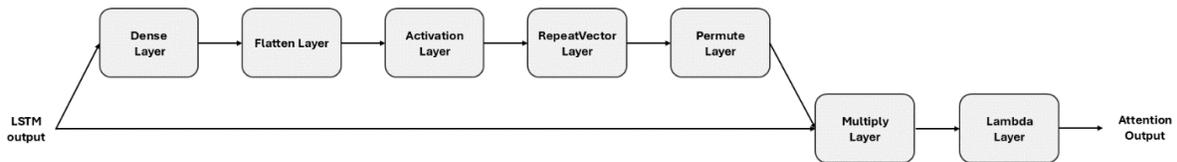


Figure 27: Custom attention mechanism structure.

The four input vectors were fed into their respective LSTM layers in a three-dimensional format: (*samples, timesteps, features*), where *samples* referred to the total number of time series samples, *timesteps* denoted the number of timesteps per time series, and *features* represented the number of features per timestep. For the CGM vector, *timesteps* was set to 288, while for the R_{ag} , R_{ai} and $Glut_4$ vectors, *timesteps* was set to 144. In all cases, the number of *features* was set to 1.

The output of each LSTM layer retained the three-dimensional format of (*samples, timesteps, units*), where units corresponded to the number of LSTM units in each layer. This output was fed into the custom attention mechanism, which produced a two-dimensional output (*samples, units*). The outputs from the four attention layers were then concatenated, resulting in a combined output of shape (*samples, units_c*), where *units_c* represented the sum of the units across the LSTM layers. The concatenated output was passed through a dropout layer to prevent overfitting, maintaining the same dimensions (*samples, units_c*). Finally, the output was fed into the final dense layer to perform the binary classification.

5.3 Model Training and Evaluation

5.3.1 Handling Missing Data and Normalizing Features

To address the issue of missing values in the CGM time series data, the KNNImputer from the scikit-learn library was utilized. This imputation method leveraged the k-nearest neighbors' algorithm to estimate and fill in missing values. Specifically, it identified the 'k' nearest samples based on the available data points and used the average of these neighbors to impute the missing values (subsection 3.6.1.1).

Various settings for the number of nearest neighbors (`n_neighbors`) were tested, including values of 3, 5, 7, and 10. After evaluating the results, the configuration with 3 nearest neighbors was chosen. This setting was selected because it best preserved the variance in the imputed data, ensuring that the imputed values closely matched the inherent variability of the original dataset.

n_neighbors	Variance	Mean
3	3488.346	159.372
5	3483.159	159.344
7	3480.818	159.347
10	3477.908	159.342

Table 4: Mean and variance statistics of the CGM vectors, following the imputation of missing values using the KNNImputer with `n_neighbors` values of 3, 5, 7, and 10.

Given that the features within each input vector operated on different scales, normalization was a crucial step to ensure that features with higher numerical values did not disproportionately influence the model compared to features with lower values. To address this issue, Z-score normalization was employed, using the StandardScaler from the scikit-learn library. The StandardScaler applied this transformation as outlined in subsection 3.6.1.2.

This method standardized each feature by transforming it according to its mean value and variance. This process resulted in features that had a mean of zero and a standard deviation of one, effectively scaling all features to have equal weight in the learning process. This normalization technique helped to balance the influence of all features, which led to a more effective and unbiased training process.

5.3.2 Hyperparameters' tuning

Throughout the network, the tanh activation function, the default in Keras, was used for the LSTM layers, and a sigmoid activation function was applied in the final dense layer to classify the samples into two categories. For training, binary cross-entropy loss function and the Adam

optimizer were used. To further optimize the network, Keras' RandomSearch was employed to experiment with different hyperparameter configurations.

Hyperparameter	Search space
Number of LSTM units	32, 64, 128
L2 Regularization factor	0.00001, 0.001, 0.1
Dropout rate	0.4, 0.6, 0.8
Learning rate	0.001, 0.01, 0.1
LSTM layer weight initializer	RandomNormal, GlorotNormal, HeNormal
Training epochs	20, 50, 100
Batch size of samples	16, 32, 64

Table 5: Parameter values checked for hyperparameter optimization.

As shown in Table 3, the OhioT1DM dataset was imbalanced, with nocturnal hypoglycemic events occurring at rates of approximately 0.33 in the training set and 0.19 in the testing set. To address this class imbalance during model training, the stratified cross-validation method was used, specifically employing the StratifiedKFold function from the scikit-learn library. Stratified k-fold cross-validation is a technique that ensures each fold has approximately the same proportion of samples from each class, making it particularly useful for imbalanced datasets.

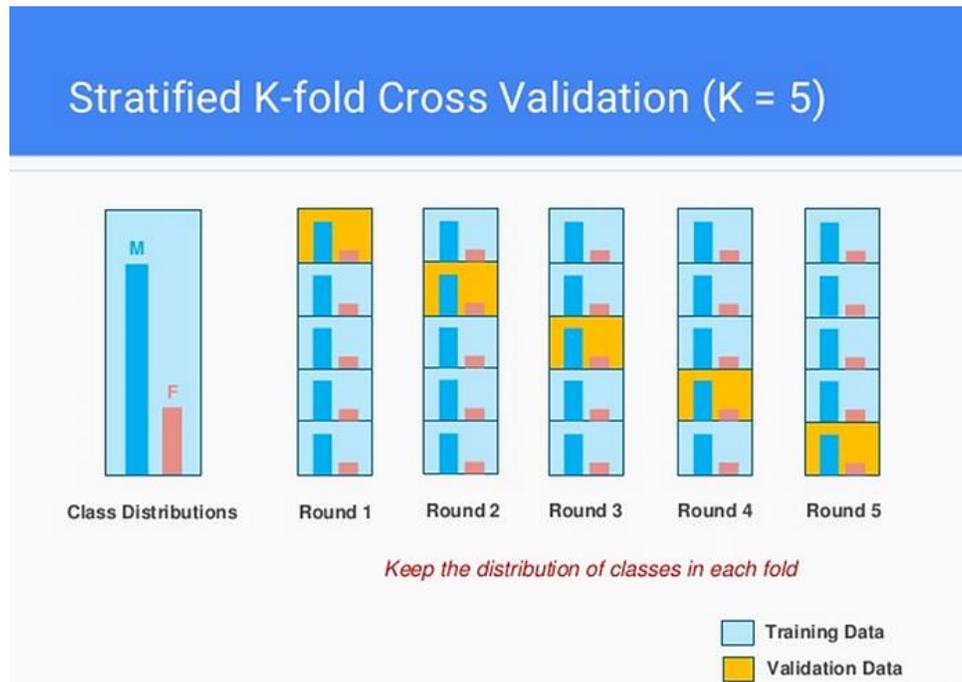


Figure 28: Example of StratifiedKFold with k=5 on an imbalanced dataset with two classes [68].

For this model, a 5-fold StratifiedKFold approach was used. This method divided the dataset into five distinct folds, each maintaining the original class distribution. In each iteration, four of the folds were combined to create the training set, while the remaining fold was used as the

validation set. This process was repeated five times, with each fold serving as the validation set exactly once. The use of StratifiedKFold ensured that the model was trained and validated on different subsets of the data, which helped in the creation of a more robust and generalizable model by preserving class balance across all folds.

5.3.3 Performance Assessment of the Model

The evaluation of the prediction model was conducted through a set of metrics designed for assessing the discrimination ability of machine learning models in binary classification tasks. A confusion matrix was used to summarize the performance, while accuracy, precision, recall, and specificity provided insights into different aspects of the model's classification abilities. Additionally, the AUC and ROC curves were utilized to assess the model's overall discriminative power.

5.3.3.1 Binary Classification Threshold Optimization

To determine the most appropriate binary classification threshold, predictions were initially generated using the training dataset. These predictions served as the basis for calculating the optimal threshold, which was identified through Youden's J statistic—a method that maximized the difference between true positive rate and false positive rate (subsection 3.6.3). Following this, the testing dataset was used for prediction, and the established threshold was applied to convert the probabilistic outputs into binary classifications. This approach ensured that the threshold was both data-driven and tailored to the specific characteristics of the model.

5.3.3.2 Prediction Uncertainty Estimation

To estimate the uncertainty of predictions, the Monte Carlo (MC) dropout technique was employed. This approach involved applying dropout during both training and inference, enabling the model to produce a distribution of predictions rather than a single output. By sampling multiple predictions for each input, the technique provided an estimate of the uncertainty associated with each prediction, which was crucial for understanding the confidence in the model's outputs [69].

5.3.3.3 External Evaluation of the Model

Finally, the proposed model underwent external evaluation using data from a ten-day monitoring period which involved 11 patients with Type 1 Diabetes, including 6 males and 5 females. These patients, who were undergoing insulin pump therapy with continuous glucose monitoring, were recruited from the Diabetes Center at the First Department of Pediatrics, P. & A. Kyriakou Children's Hospital in Athens, as part of the SMARTDIAB project. During the monitoring period, comprehensive data were collected every five minutes, including CGM readings, basal insulin levels, insulin boluses, and the amount of carbohydrates consumed in each meal [70].

Patient	Age	Gender	T1DM Duration	BMI	% HbA1c
1	3	Male	2	18.9	6.3
2	3	Male	1	16.7	8
3	9	Female	5	16.5	6.6
4	13	Female	12	24.3	7.5
5	18	Female	12	18.3	7.2
6	18	Male	10	18.6	5.7
7	18	Male	16	19.8	7.8
8	20	Female	10	24.4	6.2
9	25	Female	17	29.2	8.3
10	35	Male	22	19.8	5.7
11	38	Male	19	27.4	6

Table 6: Age in years, gender, T1DM duration in years, BMI in kg/m² and percentage of HbA1c for each patient in SMARTDIAB dataset.

Similar to the OhioT1DM study, a fixed sleep window from 8:00 PM to 8:00 AM the following morning was established, creating a 12-hour nocturnal prediction horizon. If a CGM reading dropped below 70 mg/dL during this period, the night was classified as a nocturnal hypoglycemia event (label = 1); otherwise, it was labeled as a non-hypoglycemia night (label = 0). Days with more than 24 consecutive missing CGM readings were excluded from the analysis. After these exclusions, 90 days of data remained, with a nocturnal hypoglycemia rate of 0.44.

Patient	Data Entries	Data Days	Noc Hypos	Noc Hypos Rate
1	2016	7	4	0.57
2	2592	9	2	0.22
3	2304	8	0	0
4	1728	6	3	0.5
5	2592	9	9	1
6	2304	8	6	0.75
7	2592	9	3	0.33
8	2592	9	3	0.33
9	2592	9	3	0.33
10	2304	8	2	0.25
11	2304	8	5	0.63

Table 7: Descriptive statistics of the SMARTDIAB data.

As outlined in the subsections 5.1 and 5.2, the same procedure was applied to the SMARTDIAB dataset in order to construct the four input vectors— CGM , R_{ag} , R_{ai} , and $Glut_4$ —required for

the machine learning model. These vectors were then used to assess the generalization ability of the proposed hybrid model.

6. Results and Discussion

In this chapter, the results of the proposed model are presented. The uncertainty of the model's predictions was analyzed to assess their reliability. Additionally, the performance of the developed model was compared with simpler models that used fewer input vectors. Attention weights were also employed to enhance the explainability of the model's predictions. Finally, the model trained on the OhioT1DM dataset was evaluated using the SMARTDIAB dataset to assess the model's generalization capability.

6.1 Discrimination Performance

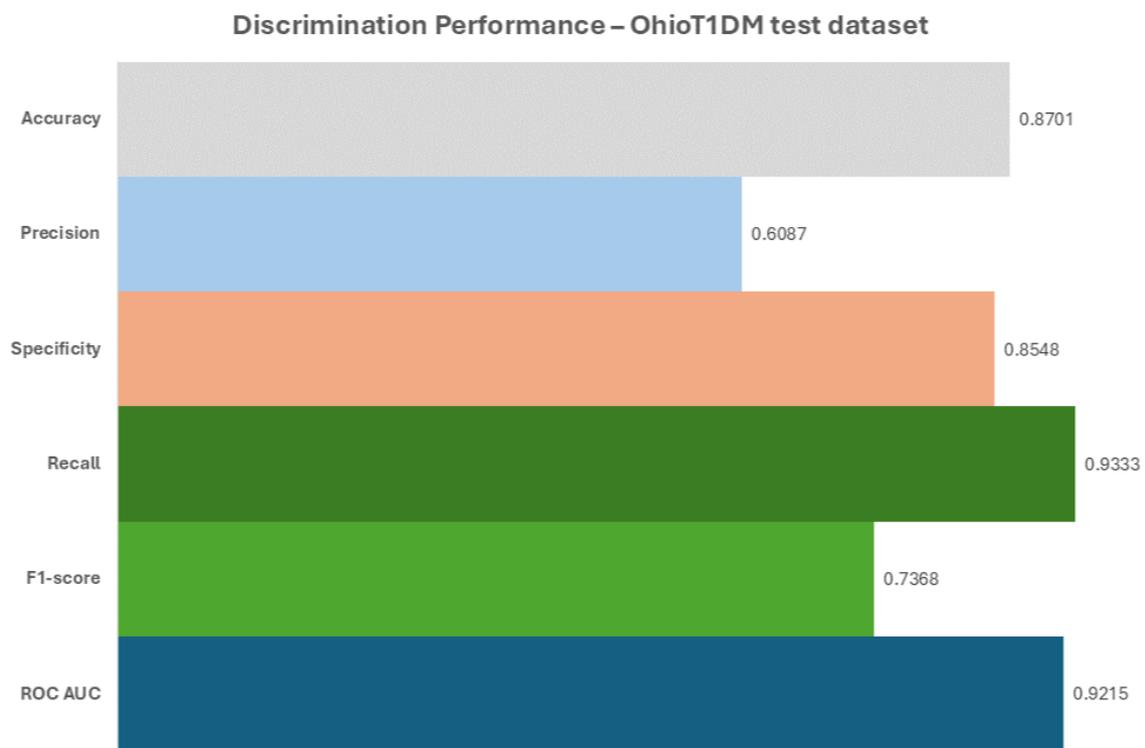


Figure 29: Evaluation results of the proposed model.

The developed model demonstrated strong performance in predicting nocturnal hypoglycemia events, achieving an accuracy of 87.01% by correctly identifying 67 out of 77 test cases. This underscored the model's overall reliability. With a precision of 60.87%, 14 out of 23 predicted nocturnal hypoglycemia events were true positives, indicating a moderate presence of false positives. The model's specificity of 85.49% highlighted its effectiveness in correctly identifying non-hypoglycemia events, reducing the likelihood of missed cases.

A recall of 93.33% showcased the model's high sensitivity, accurately detecting true nocturnal hypoglycemia events. The F1-score of 73.68% reflected a well-balanced performance between

precision and recall. Additionally, the ROC AUC score of 92.15% further confirmed the model's ability to effectively distinguish between nocturnal hypoglycemia and non-hypoglycemia events.

6.2 Comparison of different input space compositions

To assess the impact of the four input vectors on the model's performance, three simpler models were developed. The first model used only the *CGM* vector, the second combined the *CGM* and *R_{ai}* vectors, and the third incorporated *CGM*, *R_{ai}*, and *R_{ag}* vectors. The obtained performance of these models was then compared to that of the proposed model, which utilized all four input vectors (*CGM*, *R_{ag}*, *R_{ai}*, and *Glut₄*). The comparison results are summarized in the table below.

Inputs	Accuracy	Precision	Specificity	Recall	F1-score	ROC AUC
<i>CGM</i>	0.8312	0.5417	0.8226	0.8667	0.6667	0.9065
<i>CGM, R_{ai}</i>	0.8571	0.5909	0.8548	0.8667	0.7027	0.943
<i>CGM, R_{ai}, R_{ag}</i>	0.8312	0.5385	0.8065	0.9333	0.6829	0.9355
<i>CGM, R_{ai}, R_{ag}, Glut₄</i>	0.8701	0.6087	0.8548	0.9333	0.7368	0.9215

Table 8: Evaluation results of the machine learning models with the different number of input vectors.

The proposed model, using all four input vectors, achieved the highest accuracy at 87.01%, outperforming the simpler models: the *CGM*-only and *CGM* + *R_{ai}* + *R_{ag}* models achieved an accuracy score of 83.12%, while the *CGM* + *R_{ai}* model reached 85.71%. Precision improved as more vectors were added, with the full model reaching 60.87%, indicating better identification of positive cases and fewer false positives. Specificity remained high (85.48%) in models that included *R_{ai}* and *Glut₄*, suggesting these vectors enhance the accurate detection of negative cases. Models with fewer vectors showed a slight drop in specificity.

Recall was highest (93.33%) for models with three or more vectors, including the proposed model, reflecting strong sensitivity in detecting true positive cases. The F1-score was also highest for the four-vector model at 73.68%. Interestingly, the ROC AUC score was highest (94.30%) for the *CGM* + *R_{ai}* model, indicating superior discriminative ability for this specific configuration. The proposed model had a slightly lower ROC AUC of 92.15%, suggesting that while the additional vectors improved accuracy, precision, and F1-score, they slightly impacted the model's ability to distinguish between classes.

Overall, the addition of more input vectors generally improved accuracy, precision, and the F1-score, with the proposed four-vector model delivering the best overall performance, despite a small reduction in ROC AUC.

6.3 Estimation of Model's Prediction Uncertainty

To quantify uncertainty, the MC dropout technique was applied, which involved running the model through 500 stochastic passes on the test dataset. The mean probabilities from these passes were calculated to represent the predictions, while the standard deviations were used as measures of uncertainty. A lower standard deviation indicated higher confidence in the prediction, while a higher standard deviation signaled greater uncertainty.

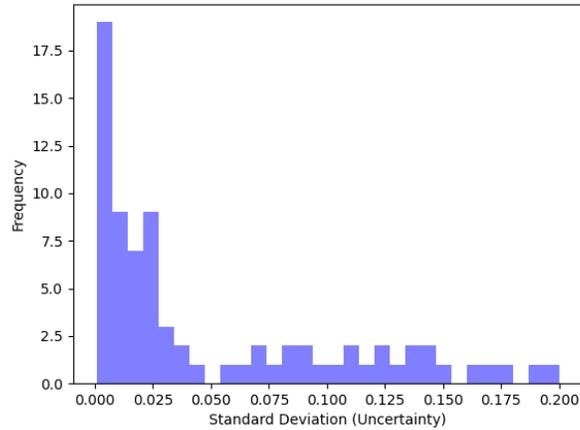


Figure 30: Distribution of Prediction Uncertainties.

The histogram above illustrates the distribution of prediction uncertainties across the OhioT1DM test dataset. Uncertainty values ranged from very low to approximately 0.2, with most predictions clustering around the lower end. This suggested that the majority of the model's predictions were made with high confidence, though a few instances showed higher uncertainty.

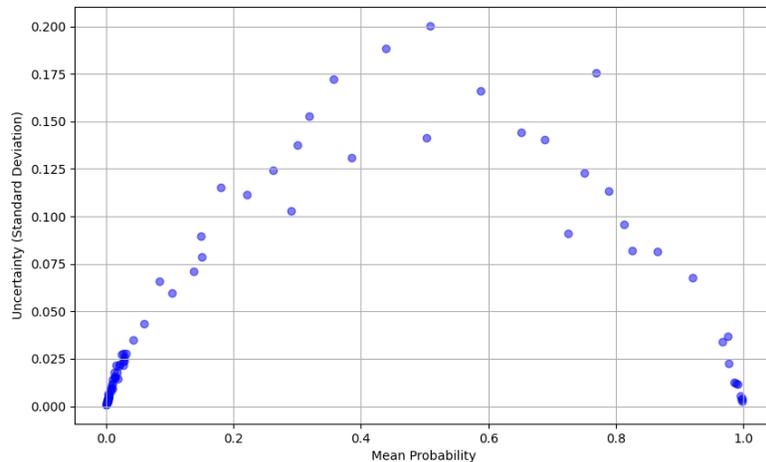


Figure 31: Uncertainty across the mean predicted probability values.

A scatter plot was generated to further explore the relationship between prediction confidence and uncertainty. It plots the average predicted probability against its corresponding uncertainty

across the 500 stochastic passes. The scatter plot shows that predictions with mean probabilities close to 1 or 0 generally have lower uncertainties, indicating strong confidence. In contrast, predictions with probabilities near the decision threshold exhibit higher uncertainties, reflecting less confidence.

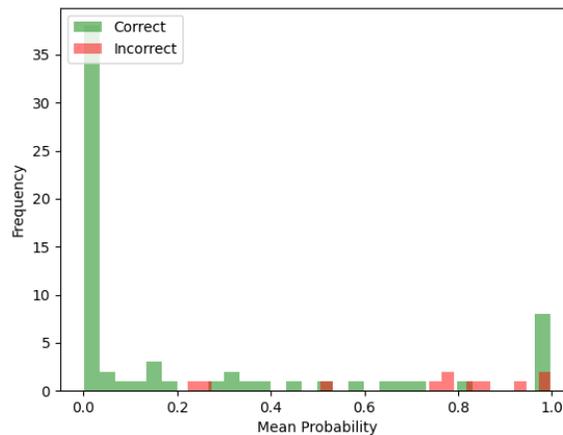


Figure 32: Overlaid histogram of mean probabilities for both correct and incorrect predictions.

To analyze the connection between confidence and accuracy, an overlaid histogram was created, comparing the mean probabilities of correct and incorrect predictions. The results reveal that correct predictions are typically associated with both high and low mean probabilities, indicating the model's strong confidence in these decisions. Incorrect predictions, however, tend to be linked with intermediate probability values, suggesting that the model was less certain and more prone to errors in these cases.

In summary, the MC dropout analysis highlights that the model's uncertainty tends to be low for predictions near the extremes, while uncertainty increases near the decision threshold, correlating with a higher likelihood of incorrect predictions.

6.4 Interpretation of Predictions Using Attention Weights

To enhance the understanding of the model's interpretability and focus across different input features, heatmaps of attention weights were generated for each input vector (CGM , R_{ag} , R_{ai} , and $Glut_4$) at each time step for all samples. Additionally, line plots were created to summarize the total attention weights across all samples for each time step.

The heatmap for the CGM input vector showed high attention scores concentrated at specific time steps across various samples, indicating that the model focused on particular periods for making predictions. The attention pattern varied across samples, suggesting that the model tailored its focus based on individual sample characteristics. The line plot revealed a decreasing trend in attention over time, with peak attention at the beginning of the previous night's sleep

and a notable increase towards the end of the previous sleep period (around the 144th time step). This pattern suggested that early glucose readings are crucial for identifying trends, while glucose levels at the end of the previous sleep period also played a significant role.

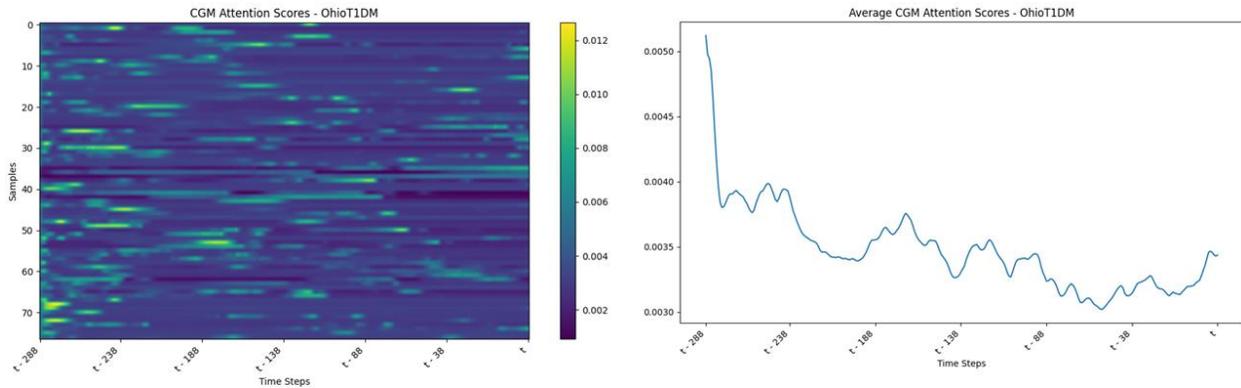


Figure 33: Visualization of attention weights for the CGM input vectors. The heatmap (left) shows attention distribution across samples and time steps, while the line plot (right) summarizes the total attention at each time step. Timestep “t” refers to the initiation of the sleep period.

The heatmap for the R_{ag} input vector indicated that attention is concentrated at specific, isolated time steps, rather than uniformly across the entire timeline. This suggested that R_{ag} had a critical impact at particular moments and samples. This observation might also be linked to missing meal information in the OhioT1DM dataset. The line plot showed a sharp increase in attention towards later time steps, especially after time step 120, reflecting the importance of glucose appearance in the bloodstream during the later sleep phases and as the awakening of the patient approaches. Early time steps had relatively low attention, highlighting that the rate of glucose appearance became more influential as time progresses.

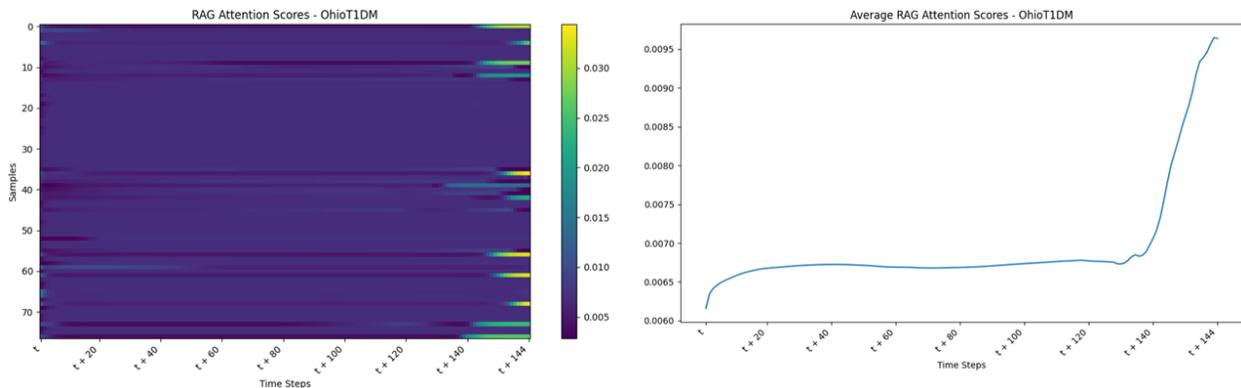


Figure 34: Visualization of attention weights for the R_{ag} input vectors. The heatmap (left) shows attention distribution across samples and time steps, while the line plot (right) summarizes the total attention at each time step. Timestep “t” refers to the initiation of the sleep period.

The heatmap for the R_{ai} input vector revealed a dispersed attention pattern, with high attention scores scattered across various time steps. This indicated that the model intermittently relied on the rate of insulin appearance throughout the timeline. The line plot demonstrated that attention

scores were initially high, drop quickly, and then rise again, peaking around the middle of the time series (around time step 50). This suggested that while the initial rate of insulin appearance was important, the model placed increasing emphasis on the middle period, likely corresponding to key metabolic changes during the sleep period (Dawn phenomenon, Somogyi effect) [71]. Attention remained relatively high in the latter part of the series, indicating its continued importance in refining predictions.

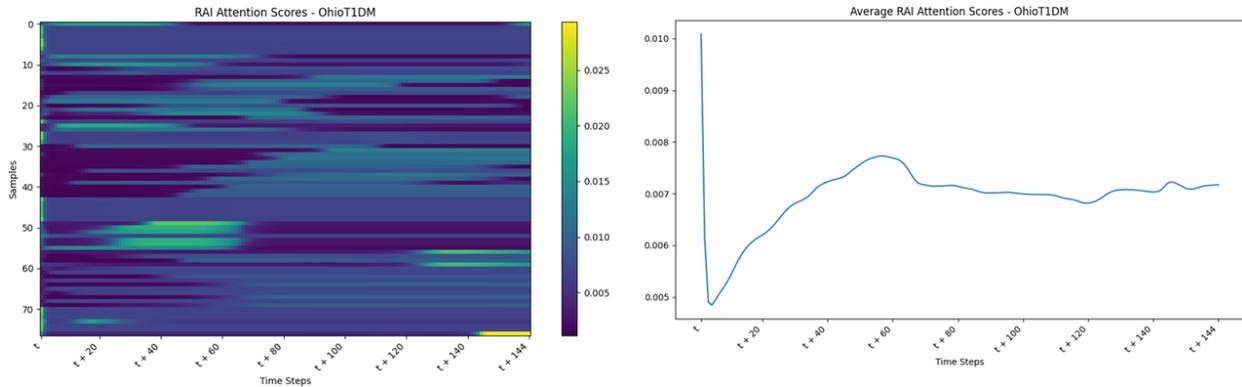


Figure 35: Visualization of attention weights for the R_{ai} input vectors. The heatmap (left) shows attention distribution across samples and time steps, while the line plot (right) summarizes the total attention at each time step. Timestep “t” refers to the initiation of the sleep period.

The heatmap for the $Glut_4$ input vector showed a varied distribution of attention, with certain time steps consistently receiving higher attention across samples. This indicated that specific moments were crucial for the model’s predictions. The line plot revealed a significant peak in attention during the early time steps, suggesting that $Glut_4$ levels at the beginning of sleep were critical for predicting glucose fluctuations. Attention decreased gradually in later time steps, reflecting the reduced impact of $Glut_4$ levels as the time sequence progresses.

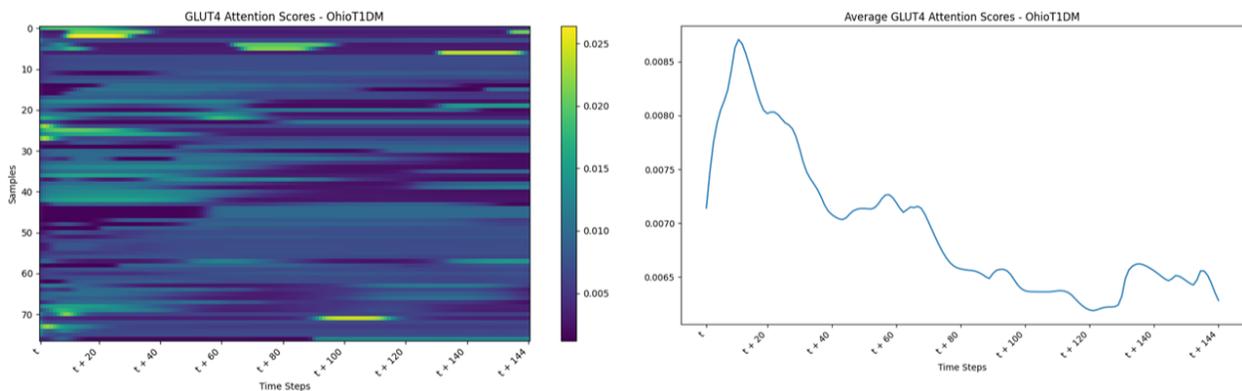


Figure 36: Visualization of attention weights for the $Glut_4$ input vectors. The heatmap (left) shows attention distribution across samples and time steps, while the line plot (right) summarizes the total attention at each time step. Timestep “t” refers to the initiation of the sleep period.

The analysis of attention weights across the four input vectors— CGM , R_{ag} , R_{ai} , and $Glut_4$ —revealed how the model dynamically prioritized different aspects of the input data over time.

The CGM input vector highlighted the importance of early and middle glucose readings. The R_{ag} vector showed that glucose appearance became more critical as the sleep period window neared. The R_{ai} vector's dispersed attention indicated intermittent reliance on insulin appearance, with significant focus during key metabolic periods. Finally, the $Glut_4$ vector emphasized the significance of early insulin-sensitive glucose receptor levels.

Overall, the model demonstrated a sophisticated mechanism for assigning varying levels of importance to different temporal features, which allowed for tailored predictions based on the most relevant information at critical moments. This detailed understanding of attention patterns not only enhanced model interpretability but also provided valuable insights into the physiological processes deemed significant for accurate hypoglycemia prediction.

6.4.1 Attention Scores Visualization for Characteristic Predictions

The attention weights of the four input vectors— CGM , R_{ag} , R_{ai} , and $Glut_4$ —were also analyzed in relation to four key prediction outcomes: true positive (TP), true negative (TN), false positive (FP), and false negative (FN). This analysis was intended to uncover how attention mechanisms vary across these prediction scenarios, providing insights into the factors contributing to both accurate and erroneous predictions.

6.4.1.1 True Positive Prediction

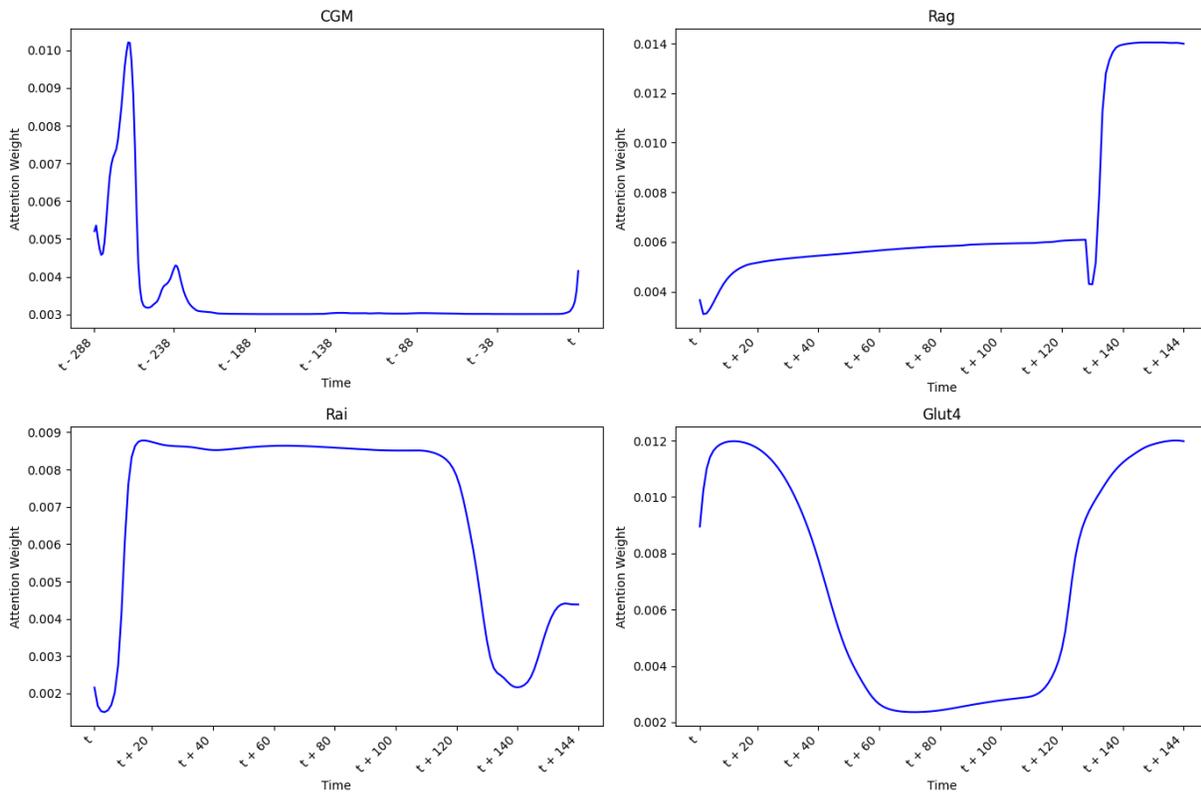


Figure 37: Attention of the four input vectors for a TP prediction. Timestep “t” refers to the initiation of the sleep period.

The attention weight for the *CGM* input vector initially started high during the early stages of the previous night's sleep, peaking early before experiencing a sharp decline. It remained low with minor fluctuations before showing a slight increase as the sleep period end approached. This pattern suggested that *CGM* values were most influential at the beginning of the previous night's sleep period, with its importance quickly diminishing and then partially recovering as the sleep period neared.

In contrast, the R_{ag} input vector began with low attention but gradually increased, stabilizing before spiking sharply towards the end of the sleep period. This trend indicated that R_{ag} significance grew over time, reaching its peak towards the conclusion of the sleep period and highlighting its crucial role in the later stages of the sleep for this specific sample.

The R_{ai} input vector demonstrated rising attention from the start, maintaining relatively high values until experiencing a sharp decline towards the end of the sleep period, followed by a slight recovery. This pattern suggested that R_{ai} was highly influential in the early stages of the sleep period but became less significant as sleep time progresses, while retaining some relevance towards the end for this particular case.

For the $Glut_4$ input vector, attention started high, dropped significantly to a low point midway through the sleep period, and then gradually increased again. This indicated that $Glut_4$ was initially important, lost significance in the middle, and regained relevance towards the end of sleep period, reflecting its varying impact at different stages of the night for this sample.

Overall, *CGM* and R_{ai} were prominent during the early stages of the time series vectors, while $Glut_4$ influence declined markedly in the middle. R_{ag} became increasingly important towards the end, with $Glut_4$ partially recovered its influence and *CGM* showed a minor uptick. This analysis underscored the dynamic role of each input vector throughout the timeline, highlighting their varying contributions at different stages for this TP prediction.

6.4.1.2 True Negative Prediction

In the TN prediction case, the *CGM* input vector exhibited sharp attention peaks early on, particularly within the first 150-time steps, which corresponded to the previous night sleep period, followed by a significant drop and consistently low attention for the remainder of the period. This indicated that *CGM* played a critical role early in the prediction process, but its influence diminished as time progresses.

The R_{ag} input started with low attention, gradually increased, and peaked near the end of the sleep period before sharply declining. This pattern suggested that R_{ag} became increasingly important as the end of the sleep period approached, particularly in the later stages. The R_{ai} input showed high attention initially, which remained elevated for the first quarter of the sleep

period before sharply declining, with a slight increase toward the end. This suggested that R_{ai} was influential early in the first stages of sleep but became less significant over time. $Glut_4$ begun with low attention, peaked early, and then gradually decreased, with some fluctuations, maintaining a more consistent role throughout the night compared to R_{ai} .

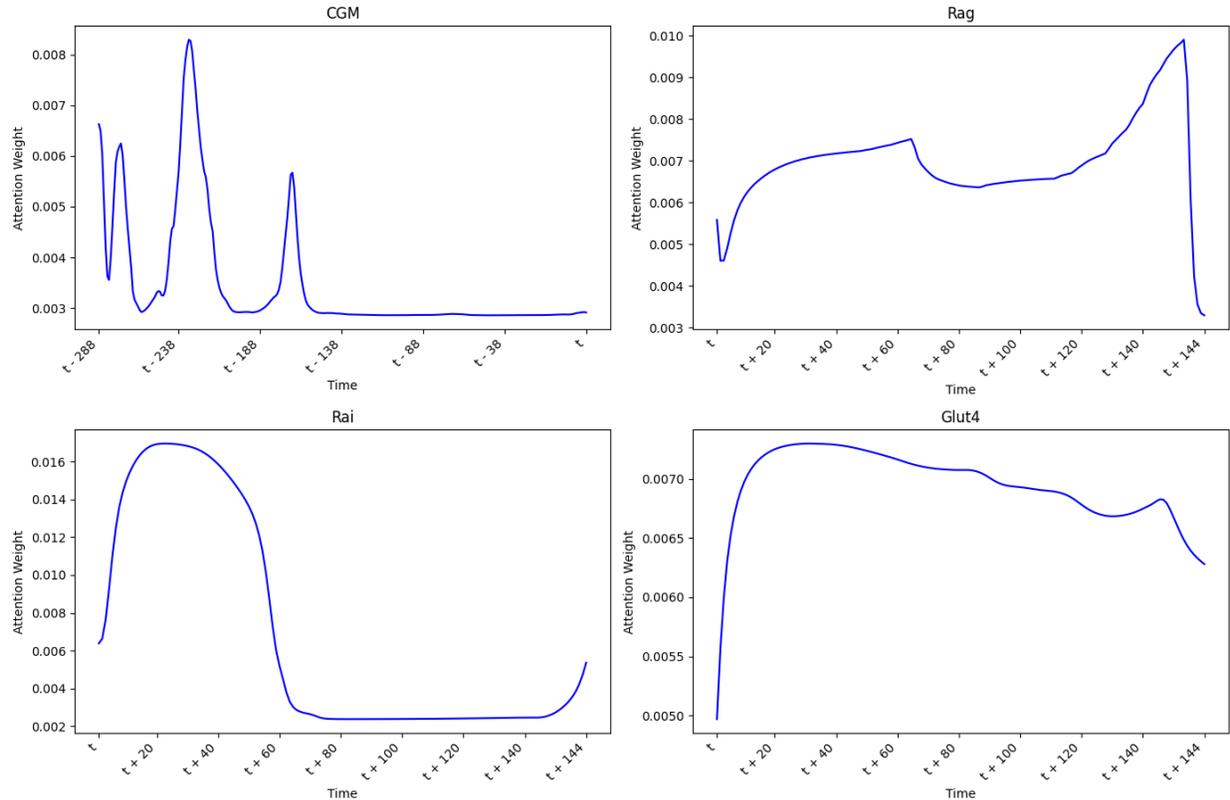


Figure 38: Attention of the four input vectors for a TN prediction. Timestep “t” refers to the initiation of the sleep period.

In summary, CGM and R_{ai} were pivotal in the early stages of TN prediction, while R_{ag} gained importance toward the end, and $Glut_4$ maintained a steadier influence. Compared to the TP prediction, which demonstrated a gradual shift in attention, TN prediction showed more abrupt changes, especially in the CGM vector. The sharp early spikes in CGM and the late rise in R_{ag} attention suggested that these inputs were key in distinguishing true negatives.

6.4.1.3 False Positive Prediction

In the FP prediction, the attention weight for the CGM input vector showed sharp spikes at the beginning, indicating that the model heavily prioritized early glucose readings of the previous night’s sleep period. After these initial peaks, attention dropped sharply and remained low for the rest of the sequence, suggesting that CGM influence was concentrated in the sleep period of the previous night. Conversely, the R_{ag} input vector gradually increased in attention, with a slight dip mid-sequence before peaking toward the end. This pattern indicated that R_{ag} became

increasingly important as the end of sleep period approached, reflecting the model's focus on the rate of glucose appearance in the bloodstream as the awakening approached.

For the R_{ai} input vector, attention remained low initially but rose sharply around time step 60, maintaining high levels for the remainder of the period. This suggested that R_{ai} became a key factor midway through the sequence and played a critical role in this FP prediction. In contrast, the $Glut_4$ input vector peaked early, followed by a gradual decline, indicating that $Glut_4$ was initially significant but diminished in importance as the sleep progressed.

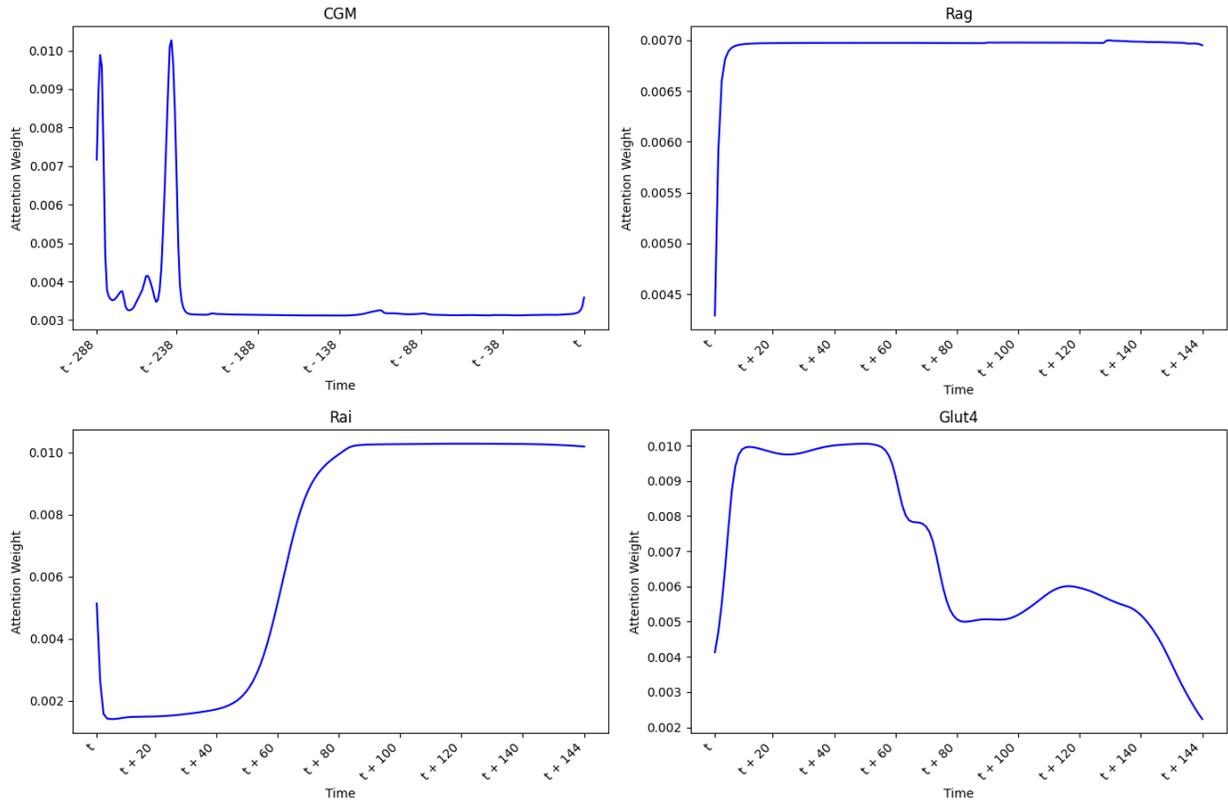


Figure 39: Attention of the four input vectors for a FP prediction. Timestep “t” refers to the initiation of the sleep period.

Overall, in this FP scenario, CGM played a strong role at the beginning, while R_{ag} and R_{ai} gained prominence later in the sequence. $Glut_4$ exerted strong early influence but gradually faded. While R_{ag} and R_{ai} also became more influential later in TP and TN predictions, this shift was more pronounced in FP scenarios, suggesting that the model might overemphasized late-stage dynamics (R_{ag} and R_{ai}) in false positives. This could led to an overestimation of glucose appearance and insulin activity, resulting in incorrect predictions.

6.4.1.4 False Negative Prediction

In the FN prediction, the attention weights for the CGM input vector exhibited sharp early peaks and significant fluctuations until time step 200, indicating intermittent focus on glucose readings.

Afterward, CGM attention decreased and remained low, reflecting reduced influence in later stages. R_{ag} , on the other hand, steadily increased in attention, peaking toward the end, indicating its growing importance as the awakening neared.

R_{ai} started with high attention but dropped sharply after time step 60, losing relevance in later stages. This decline suggested the model initially valued insulin-related signals but failed to account for their ongoing impact, which might led to missed cues in the prediction. $Glut_4$, in contrast, begun with low attention but gradually rose, becoming more influential toward the end of the sequence. This increase indicated the model's growing reliance on glucose regulation factors, though this late-stage focus on $Glut_4$ might not fully counterbalanced the early decline in R_{ai} , contributing to the false negative prediction.

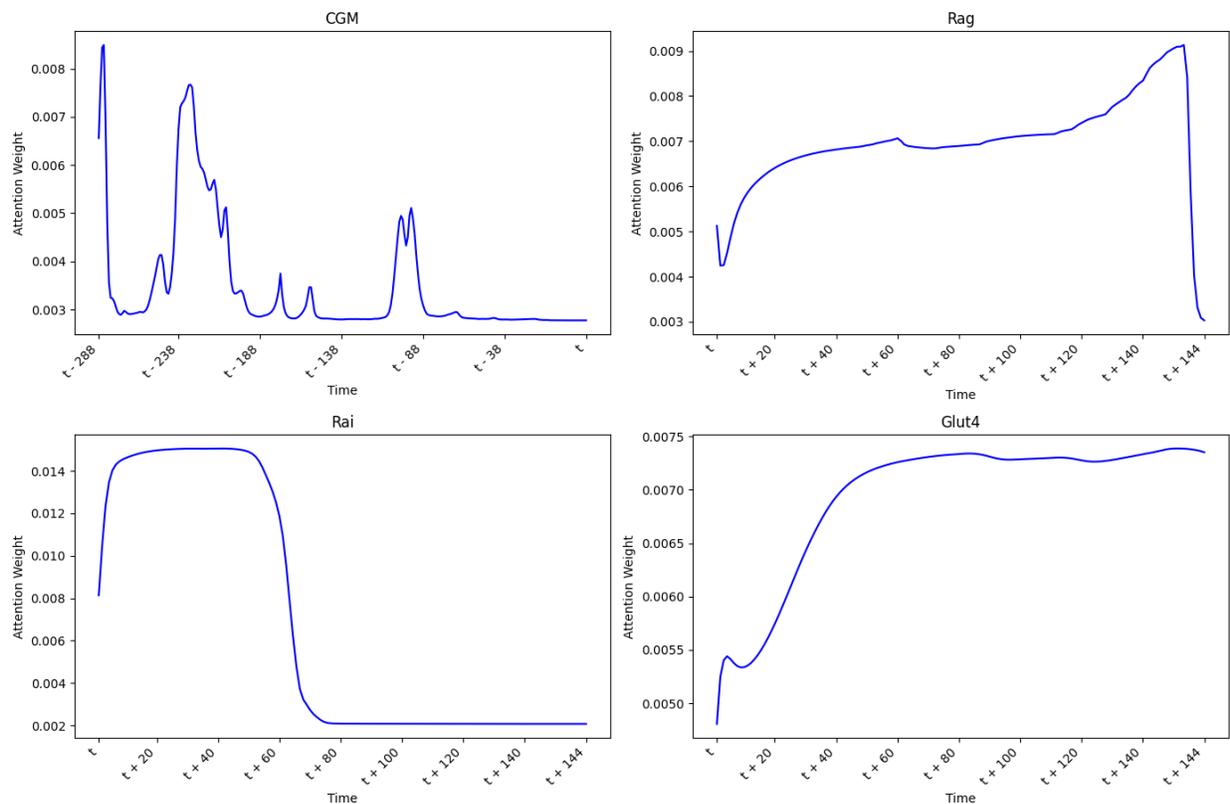


Figure 40: Attention of the four input vectors for a FN prediction. Timestep “t” refers to the initiation of the sleep period.

In this FN prediction, attention patterns—particularly for CGM —showed more variability and inconsistency compared to the stable patterns seen in TP and TN predictions. The erratic focus on CGM and the sharp early decline in R_{ai} importance suggested that critical information was overlooked, leading to the prediction error. While R_{ag} and $Glut_4$ exhibited steady increases, similar to their roles in TP and TN predictions, the inconsistent weighting of CGM and the early drop in R_{ai} relevance likely played a major role in the model's failure to accurately predict the outcome.

6.5 Evaluation of the Model on the SMARTDIAB dataset

6.5.1 Discrimination Performance

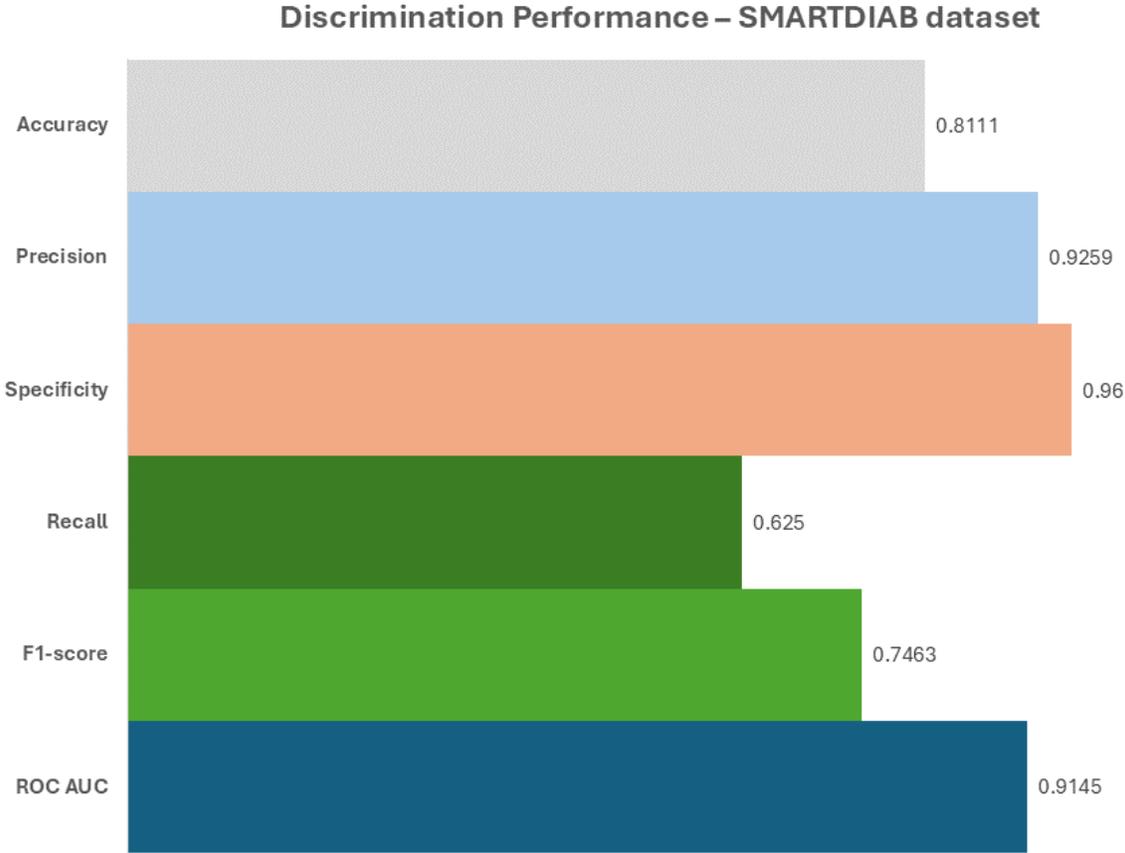


Figure 41: Evaluation results of the machine learning model in the SMARTDIAB dataset.

The model was evaluated on the SMARTDIAB dataset, yielding an accuracy of 81.11%. This indicated that it correctly predicted 73 out of 90 samples, a slightly lower rate than in the OhioT1DM test dataset (Table 8), but still demonstrating strong overall performance. The precision of 92.59% showed a significant improvement, with 25 out of 27 predicted nocturnal hypoglycemia events being true positives, suggesting that the model had effectively reduced the occurrence of false positives in this dataset’s predictions.

The model's specificity of 96.00% was particularly notable, as it demonstrated a heightened ability to correctly identify non-hypoglycemia cases. However, the recall dropped to 62.50%, indicating that while the model remained highly precise, it missed some true hypoglycemia cases. The F1-score of 74.63% reflected a strong balance between precision and recall. Despite this, the ROC AUC score of 91.45% underscored the model's continued performance in distinguishing

between nocturnal hypoglycemia and non-hypoglycemia events, confirming its robustness across different datasets.

In summary, the model continued to perform effectively in predicting nocturnal hypoglycemia, with a marked improvement in precision and specificity, though it faced some challenges in maintaining high recall.

6.5.2 Estimation of Model’s Prediction Uncertainty

The uncertainty in the model's predictions on the new dataset was evaluated using the same MC dropout technique as applied in the original analysis. The model was run over 500 stochastic passes, and the mean probabilities of these passes were calculated to assess prediction confidence, while the corresponding standard deviations were used as a measure of uncertainty. As in the previous evaluation, a lower standard deviation reflected greater prediction confidence, while a higher standard deviation suggested higher uncertainty.

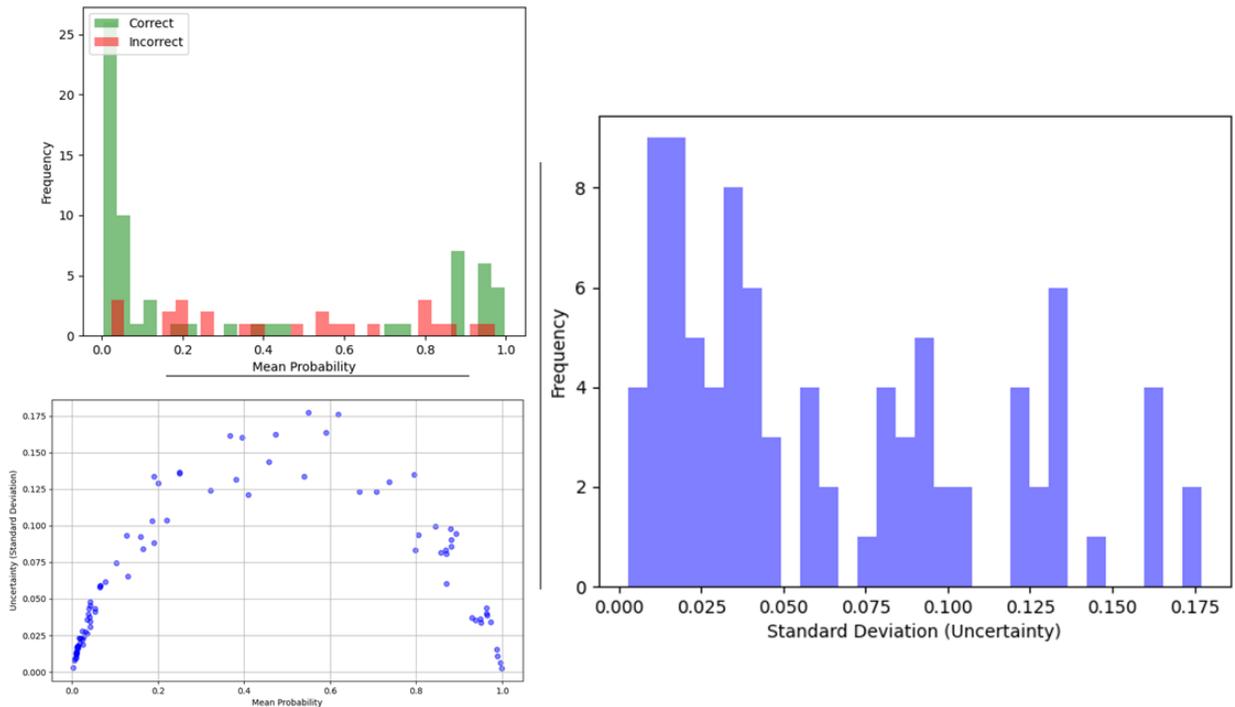


Figure 42: Visualization of predictions’ uncertainty in the SMARTDIAB dataset. Top left: Overlaid histogram of mean probabilities for both correct and incorrect predictions. Bottom left: Uncertainty across the mean predicted probability values. Right: Distribution of Prediction Uncertainties

The uncertainty histogram for the SMARTDIAB dataset closely mirrored that of the previous evaluation, with most uncertainty values concentrated at the lower end, signifying that the majority of predictions were made with high confidence. A small number of predictions, however, showed higher uncertainty, with values reaching around 0.175, indicating less certainty in some cases.

The scatter plot mapping averaged probabilities against their uncertainties revealed a pattern similar to that seen in the OhioT1DM dataset. Predictions with mean probabilities near 1 or 0 had lower associated uncertainties, reflecting the model's confidence in these cases. In contrast, predictions around the decision threshold exhibited higher uncertainty, highlighting the model's uncertainty in these borderline instances.

The overlaid histogram comparing correct and incorrect predictions showed a similar pattern. Correct predictions tended to cluster near the extremes, close to 0 or 1, reflecting high confidence. Incorrect predictions, by contrast, were more frequently associated with intermediate probability values, where the model showed lower confidence and higher uncertainty.

6.5.3 Interpretation of Predictions Using Attention Weights

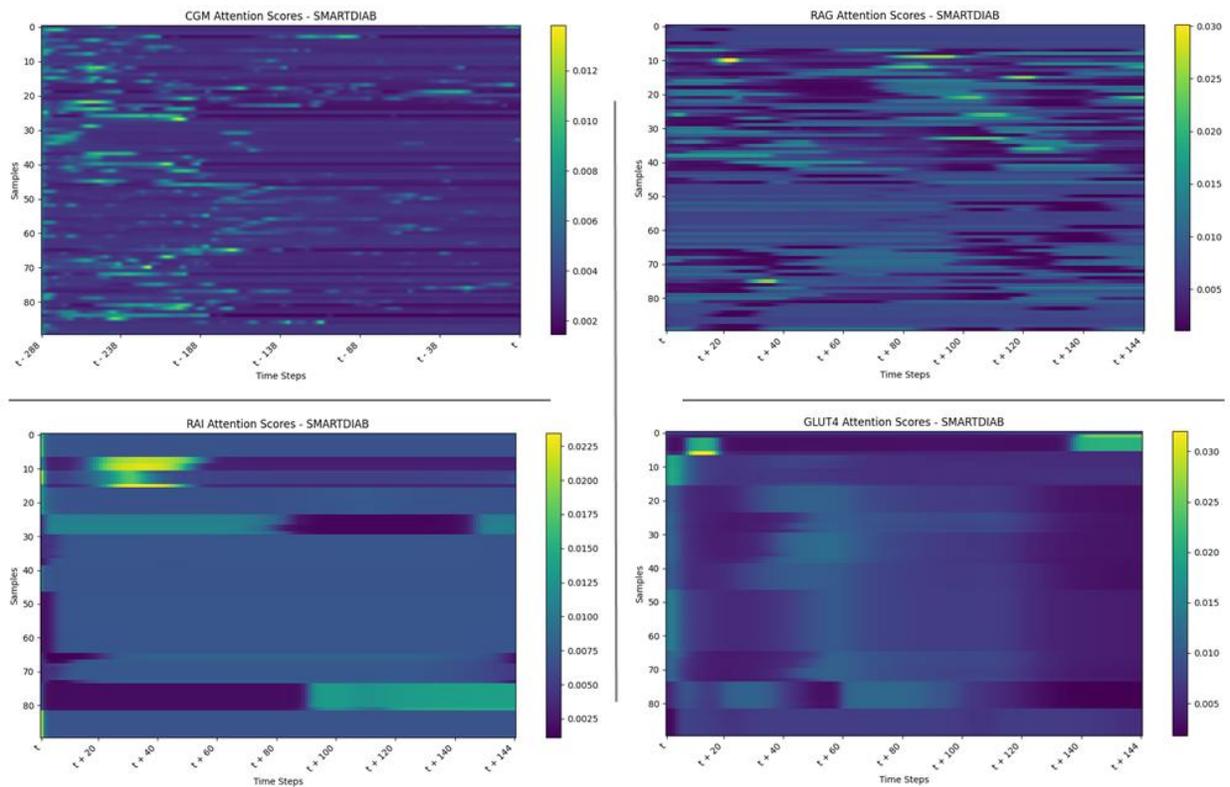


Figure 43: Heatmaps of the attention weights of the four input vectors of the SMARTDIAB dataset. Timestep “t” refers to the initiation of the sleep period.

The *CGM* input vector's heatmap showed that the model concentrated on specific time points, though the attention was more diffuse than in the OhioT1DM dataset. This suggested that in the SMARTDIAB dataset, the model spread its reliance on *CGM* data across the timeline, likely due to variability in glucose patterns among the samples. For the R_{ag} input vector, the heatmap indicated that the model selectively focused on a few critical moments in the sequence, reflecting the importance of glucose appearance at particular time points.

The R_{ai} heatmap showed a more dispersed and sporadic attention pattern, signaling variable insulin dynamics across samples, with certain periods attracting more attention depending on sample-specific factors. The $Glut_4$ heatmap similarly revealed variability across samples and time steps, with certain moments consistently drawing more focus, while much of the sequence showed lower attention. This suggested a more intermittent influence of $Glut_4$ on the model's predictions in this dataset.

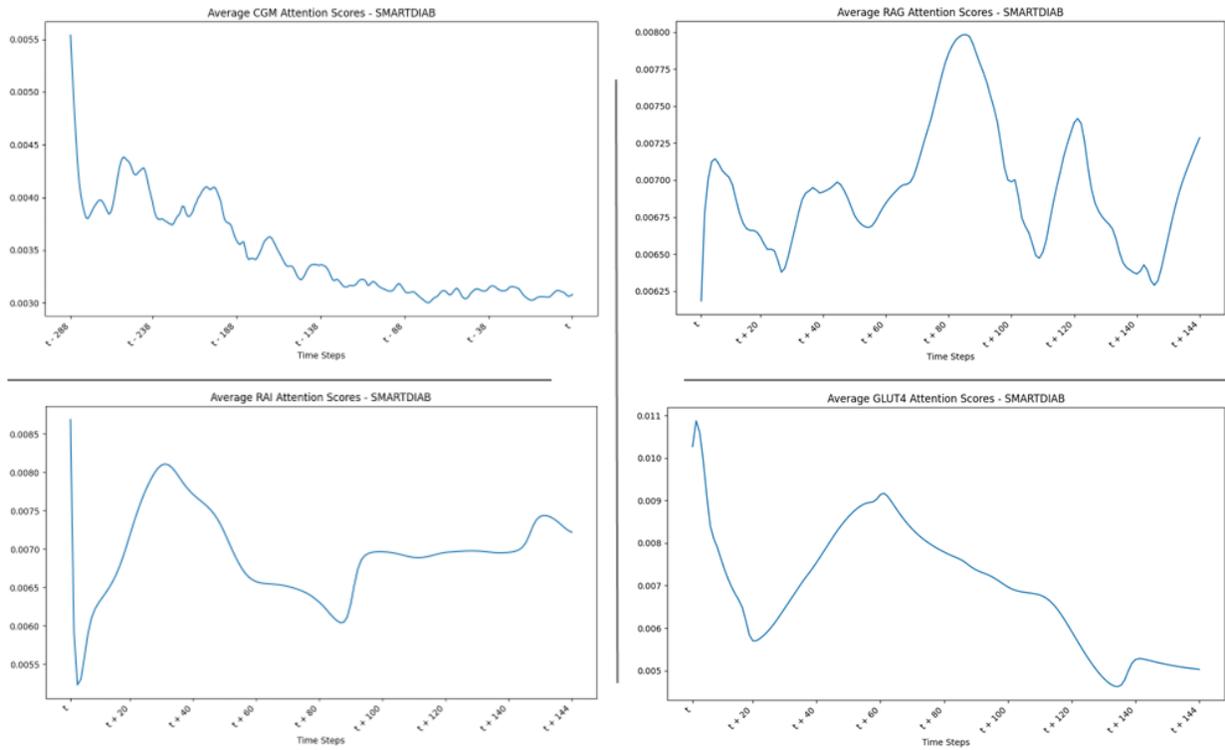


Figure 44: Line plots of the averaged attention weights of the four input vectors of the SMARTDIAB dataset. Timestep “t” refers to the initiation of the sleep period.

In the line plot for CGM , a clear trend emerged: the model allocated higher attention early in the time sequence, which gradually decreased as time progresses. This highlighted the importance of early glucose readings from the previous night’s sleep period in identifying trends. Lower attention in later time steps suggested that the model relied less on these later glucose readings compared to the previous dataset. The R_{ag} line plot showed a significant increase in attention toward the middle and later time steps, emphasizing the role of glucose appearance after fasting in shaping glucose levels during the sleep period. This shifting attention aligned with the physiological importance of glucose dynamics during these periods, as the end of sleep period approached.

For the R_{ai} input vector, attention rose during the middle time steps, although less sharply than in the previous dataset. Higher attention early and midway through the timeline indicated the model's continued reliance on insulin appearance during early and middle sleep periods. This

pattern underscored the role of insulin data in refining predictions. Finally, the $Glut_4$ line plot showed a prominent early peak, indicating the model prioritized initial $Glut_4$ levels. This reflected the critical role of insulin-sensitive glucose receptors at the start of the night, though attention declined toward later time steps, suggesting early $Glut_4$ levels were more pivotal in this dataset.

Overall, the SMARTDIAB dataset analysis revealed that the model continued to prioritize different features dynamically across time. Early glucose readings remained crucial, while R_{ag} became more important in later time steps, reflecting the physiological significance of glucose appearance during fasting sleeping periods. The R_{ai} vector showed intermittent attention, particularly in the middle of the sequence, highlighting its role in insulin dynamics. $Glut_4$ remained influential in the early stages, reinforcing its impact on shaping glucose trends at the beginning of the night.

6.6 Comparison with the State of the Art

Study	Dataset	Model	AUC-ROC	Specificity	Recall
Berikov et al. [61]	406 T1D patients (clinic database)	RF	0.94	0.87	0.90
Jensen et al. [58]	463 T1D patients (Onset 5 trial by Novo Nordisk)	LDA	0.79	0.70	0.75
Parcerisas et al. [60]	10 T1D patients (clinic database)	SVM	0.81	0.75	0.95
Mosquera-Lopez et al. [59]	124 T1D patients (Tidepool Big Data Donation Dataset)	SVR	0.86	0.94	0.72
Guemes et al. [72]	6 T1D patients (OhioT1DM dataset)	SVM	0.78	0.72	0.68
Afentakis et al. [73]	37 T1D patients (clinic database)	SVM	0.79	0.72	0.72
Proposed model	12 T1D patients (OhioT1DM dataset)	LSTM	0.92	0.85	0.93

Table 9: Comparison of model performance metrics in T1D nocturnal hypoglycemia prediction studies. RF: Random Forest, LDA: Linear Discrimination Analysis, SVM: Support Vector Machine, SVR: Support Vector Regression.

In comparison to existing studies, the developed model demonstrated competitive performance across several key metrics. With an AUC-ROC of 0.9215, it performed similarly to Berikov’s model and outperformed other models like Jensen’s and Parcerisas’. Additionally, the developed model’s specificity (0.8548) was close to top performers like Mosquera-Lopez’s model, indicating strong ability to correctly identify negative cases.

Despite the lower precision, the proposed model yielded a recall score of 0.9333, which surpassed most compared models, ensuring that the majority of positive cases were accurately

identified. Overall, the proposed model demonstrated satisfying discrimination performance, with particularly strong recall and AUC-ROC values.

7. Conclusions and Future Work

In this study, an interpretable model was developed to estimate the risk of nocturnal hypoglycemia in patients with Type 1 Diabetes. The model was trained and evaluated using real patient data from the OhioT1DM dataset, which included CGM, insulin administration, and carbohydrate intake. To capture the complex physiological processes of insulin absorption and glucose metabolism, simulation models were applied to transform the input data. These models calculated physiological parameters such as the rates of glucose and insulin appearance in the blood and the percentage of GLUT4 transporters on cell membranes during sleep.

The predictions were generated using an RNN architecture that incorporated one LSTM layer and an Attention layer for each input vector. A Concatenate layer, Dropout layer, and Dense layer followed to complete the model architecture. Hyperparameter tuning was performed using stratified cross-validation, and the binary classification threshold was determined by Youden's J Statistic. Importantly, the model's interpretability was enhanced by visualizing the Attention layer weights for each input vector.

The model exhibited strong performance, achieving an accuracy of 87.01% and a recall of 93.33%, demonstrating its ability to detect nocturnal hypoglycemia events effectively. The F1-score of 73.68% and precision of 60.87% suggested solid overall performance, albeit with moderate false-positive rates. A ROC AUC of 92.15% underscored the model's effectiveness in distinguishing between nocturnal hypoglycemia and non-hypoglycemia events.

Analysis of the attention weights provided key insights into the model's decision-making process. The model dynamically prioritized input vectors such as blood glucose levels (*CGM*), rates of glucose and insulin appearance (R_{ag} and R_{ai}), and GLUT4 levels ($Glut_4$) based on the time steps and physiological changes during sleep. *CGM* data from the previous night's sleep period and $Glut_4$ percentage on the cell surface at sleep onset were more influential. As the sleep progressed, R_{ag} gain importance, particularly near the "wake-up time", while R_{ai} showed intermittent relevance, reflecting how the model adapts to metabolic shifts during the fasting state of sleep. Furthermore, MC dropout analysis revealed that predictions with probabilities near the extremes had low uncertainty, while those close to the decision threshold exhibited higher uncertainty, often correlating with prediction errors. This indicated the model's confidence in its predictions and highlighted areas for improvement in handling borderline cases.

The model's robustness was further tested on the SMARTDIAB dataset, where it maintained high specificity (96%) and improved precision (92.59%). However, recall (62.50%) dropped, indicating the model occasionally missed true hypoglycemia cases. This suggested that while the model performs well in identifying positive cases, it may struggle with certain patterns in different datasets.

Several potential future improvements can be outlined here. First, the parameterization of the simulation models could be adjusted by incorporating real data on sleep onset and duration, rather than relying on hypothesized sleep windows. Additionally, the fixed 12-hour sleep horizon could be reconsidered. Since the model is global, incorporating patient-specific data such as age, weight, and gender might enhance personalization. With sufficient data, fully individualized models could be developed.

In terms of model architecture, experimenting with alternative designs could yield better results. Although this study only explored one architecture, refining the decision threshold, enhancing the attention mechanism, or employing post-processing techniques could reduce false positives and negatives. The decline in recall on the SMARTDIAB dataset underscores the need to improve recall without compromising precision, perhaps through ensemble methods or dataset-specific feature engineering [74, 75].

Further improvements might come from incorporating additional physiological or behavioral inputs such as physical activity, illness, or sleep quality, which could enhance the model's accuracy. Exploring physiological biomarkers beyond *CGM*, R_{ag} , R_{ai} , and $Glut_4$ could provide a more holistic understanding of nocturnal hypoglycemia patterns. Moreover, addressing prediction uncertainty through methods like uncertainty-aware learning or confidence calibration could improve the reliability of borderline predictions [76].

Finally, interpretability could be enhanced through techniques such as attention maps, SHAP SHapley Additive exPlanations (SHAP), or Local Interpretable Model-agnostic Explanations (LIME), making the model more transparent and trustworthy in clinical applications [77, 78]. Further investigation into attention mechanisms might also offer deeper insights into the physiological factors driving hypoglycemia during sleep.

8. References

- [1] L. Stryer, *Biochemistry, 4th Edition*. New York: W. H. Freeman and Company, 1995.
- [2] S. L. Aronoff, K. Berkowitz, B. Shreiner, and L. Want, "Glucose Metabolism and Regulation: Beyond Insulin and Glucagon," *Diabetes Spectrum*, vol. 17, pp. 183-190, 2004.
- [3] D. Singh, A. Shati, M. Alfaifi, S. Elbehairi, I. Han, and E.-H. Choi, "Development of Dementia in Type 2 Diabetes Patients: Mechanisms of Insulin Resistance and Antidiabetic Drug Development," *Cells*, vol. 11, p. 3767, 11/25 2022.
- [4] P. Sonksen and J. Sonksen, "Insulin: understanding its action in health and disease," *British Journal of Anaesthesia*, vol. 85, pp. 69-79, 2000.
- [5] M. S. Rahman, K. S. Hossain, S. Das, S. Kundu, E. O. Adegoke, M. A. Rahman, *et al.*, "Role of Insulin in Health and Disease: An Update," *International Journal of Molecular Sciences*, vol. 22, p. 6403, 2021.
- [6] G. Jiang and B. B. Zhang, "Glucagon and regulation of glucose metabolism," *American Journal of Physiology-Endocrinology and Metabolism*, vol. 284, pp. E671-E678, 2003.
- [7] W. H. Organization. (2023). *Diabetes*. Available: <https://www.who.int/news-room/fact-sheets/detail/diabetes>
- [8] R. Barbie Cervoni. (2024). *Is Diabetes Genetic?*
Genetic Risks, Testing, and Prevention. Available: <https://www.verywellhealth.com/is-diabetes-genetic-5112506>
- [9] A. D. Association, "2. Classification and Diagnosis of Diabetes: Standards of Medical Care in Diabetes—2021," *Diabetes Care*, vol. 44, pp. S15-S33, 2020.
- [10] P. J. Kumar and M. L. Clark, *Kumar & Clark clinical medicine*, 5th ed ed. Edinburgh: Saunders, 2002.
- [11] A. Rewers, "Acute Metabolic Complications in Diabetes," in *Diabetes in America*, C. C. Cowie, S. S. Casagrande, A. Menke, M. A. Cissell, M. S. Eberhardt, J. B. Meigs, *et al.*, Eds., ed Bethesda (MD): National Institute of Diabetes and Digestive and Kidney Diseases (US), 2018.
- [12] G. Umpierrez and M. Korytkowski, "Diabetic emergencies — ketoacidosis, hyperglycaemic hyperosmolar state and hypoglycaemia," *Nature Reviews Endocrinology*, vol. 12, pp. 222-232, 2016/04/01 2016.
- [13] A. J. Graveling and B. M. Frier, "The risks of nocturnal hypoglycaemia in insulin-treated diabetes," *Diabetes Res Clin Pract*, vol. 133, pp. 30-39, Nov 2017.
- [14] K. V. Allen and B. M. Frier, "Nocturnal Hypoglycemia: Clinical Manifestations and Therapeutic Strategies Toward Prevention," *Endocrine Practice*, vol. 9, pp. 530-543, 2003/11/01/ 2003.
- [15] D. M. Nathan, "Long-term complications of diabetes mellitus," *N Engl J Med*, vol. 328, pp. 1676-85, Jun 10 1993.
- [16] D. P. K. Uppula. (2024). *What are the Long term Complications of Diabetes?* Available: <https://diabetes.co.in/what-are-the-long-term-complications-of-diabetes/>
- [17] N. E. Institute. (2024). *Diabetic Retinopathy*. Available: <https://www.nei.nih.gov/learn-about-eye-health/eye-conditions-and-diseases/diabetic-retinopathy#section-id-10545>
- [18] A. Lim, "Diabetic nephropathy - complications and treatment," *Int J Nephrol Renovasc Dis*, vol. 7, pp. 361-81, 2014.
- [19] B. M. Leon and T. M. Maddox, "Diabetes and cardiovascular disease: Epidemiology, biological mechanisms, treatment recommendations and future research," *World J Diabetes*, vol. 6, pp. 1246-58, Oct 10 2015.

- [20] A. J. M. Boulton, A. I. Vinik, J. C. Arezzo, V. Bril, E. L. Feldman, R. Freeman, *et al.*, "Diabetic Neuropathies: A statement by the American Diabetes Association," *Diabetes Care*, vol. 28, pp. 956-962, 2005.
- [21] A. D. Association, "Diagnosis and classification of diabetes mellitus," *Diabetes Care*, vol. 33 Suppl 1, pp. S62-9, Jan 2010.
- [22] A. American Diabetes, "Standards of Medical Care in Diabetes—2022 Abridged for Primary Care Providers," *Clinical Diabetes*, vol. 40, pp. 10-38, 2022.
- [23] K. Zarkogianni, E. Litsa, K. Mitsis, P.-Y. Wu, C. Kaddi, C. Cheng, *et al.*, "A Review of Emerging Technologies for the Management of Diabetes Mellitus," *IEEE transactions on bio-medical engineering*, vol. 62, 08/19 2015.
- [24] K. Zarkogianni and K. Nikita, "Personal Health Systems for Diabetes Management, Early Diagnosis and Prevention," ed, 2015, pp. 465-492.
- [25] Z. M. Mathew TK, Tadi P. (2023). *Blood Glucose Monitoring*. Available: <https://www.ncbi.nlm.nih.gov/books/NBK555976/>
- [26] G. Cappon, M. Vettoretti, G. Sparacino, and A. Facchinetti, "Continuous Glucose Monitoring Sensors for Diabetes Management: A Review of Technologies and Applications," *Diabetes Metab J*, vol. 43, pp. 383-397, Aug 2019.
- [27] A. Soni, N. Wright, J. C. Agwu, A. Timmis, J. Drew, M. Kershaw, *et al.*, "A practical approach to continuous glucose monitoring (rtCGM) and FreeStyle Libre systems (isCGM) in children and young people with Type 1 diabetes," *Diabetes Research and Clinical Practice*, vol. 184, 2022.
- [28] A. D. Association, "Insulin Administration," *Diabetes Care*, vol. 27, pp. s106-s107, 2004.
- [29] C. Berget, L. H. Messer, and G. P. Forlenza, "A Clinical Overview of Insulin Pump Therapy for the Management of Diabetes: Past, Present, and Future of Intensive Therapy," *Diabetes Spectrum*, vol. 32, pp. 194-204, 2019.
- [30] A. D. Association, "Nutrition Recommendations and Interventions for Diabetes: A position statement of the American Diabetes Association," *Diabetes Care*, vol. 31, pp. S61-S78, 2008.
- [31] S. R. Colberg, R. J. Sigal, J. E. Yardley, M. C. Riddell, D. W. Dunstan, P. C. Dempsey, *et al.*, "Physical Activity/Exercise and Diabetes: A Position Statement of the American Diabetes Association," *Diabetes Care*, vol. 39, pp. 2065-2079, 2016.
- [32] C. C. Aggarwal, *Neural Networks and Deep Learning*: Springer Cham, 2023.
- [33] S. Haykin, *Neural Networks and Learning Machines, Third Edition*. Upper Saddle River, New Jersey 07458: Pearson Education, Inc., 2009.
- [34] S. Sharma, S. Sharma, and A. Athaiya, "ACTIVATION FUNCTIONS IN NEURAL NETWORKS," *International Journal of Engineering Applied Sciences and Technology*, vol. 04, pp. 310-316, 05/10 2020.
- [35] K. Y. Chan, B. Abu-Salih, R. Qaddoura, A. M. Al-Zoubi, V. Palade, D.-S. Pham, *et al.*, "Deep neural networks in the cloud: Review, applications, challenges and research directions," *Neurocomputing*, vol. 545, p. 126327, 2023/08/07/ 2023.
- [36] Z. C. Lipton, "A Critical Review of Recurrent Neural Networks for Sequence Learning," *CoRR*, vol. abs/1506.00019, 2015.
- [37] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, pp. 1735-1780, 1997.
- [38] C. Olah, "Understanding LSTM Networks," in *colah's blog*, ed, 2015.
- [39] N. Arya, "Three Types of Machine Learning," in *EJable*, ed, 2024.
- [40] M. W. Heymans and J. W. R. Twisk, "Handling missing data in clinical research," *Journal of Clinical Epidemiology*, vol. 151, pp. 185-188, 2022.
- [41] Simplilearn. (2023). *Introduction to Data Imputation*. Available: <https://www.simplilearn.com/data-imputation-article>

- [42] D. Singh and B. Singh, "Investigating the impact of data normalization on classification performance," *Applied Soft Computing*, p. 105524, 05/01 2019.
- [43] S. Sontakke, J. Lohokare, R. Dani, and P. Shivagaje, "Classification of Cardiotocography Signals Using Machine Learning: Proceedings of the 2018 Intelligent Systems Conference (IntelliSys) Volume 2," ed, 2019, pp. 439-450.
- [44] D. W. Gareth James, Trevor Hastie, Robert Tibshirani, *An Introduction to Statistical Learning with Applications in R*: Springer New York, NY, 2021.
- [45] M. Hossin and S. M.N, "A Review on Evaluation Metrics for Data Classification Evaluations," *International Journal of Data Mining & Knowledge Management Process*, vol. 5, pp. 01-11, 03/31 2015.
- [46] M. Rodríguez-Hernández, R. Pruneda, and J. Rodríguez-Díaz, "Statistical Analysis of the Evolutive Effects of Language Development in the Resolution of Mathematical Problems in Primary School Education," *Mathematics*, vol. 9, p. 1081, 05/11 2021.
- [47] W. J. Youden, "Index for rating diagnostic tests," *Cancer*, vol. 3, pp. 32-5, Jan 1950.
- [48] S. Oviedo, J. Vehí, R. Calm, and J. Armengol, "A review of personalized blood glucose prediction strategies for T1DM patients," *International journal for numerical methods in biomedical engineering*, vol. 33, 09/19 2016.
- [49] R. N. Bergman, Y. Z. Ider, C. R. Bowden, and C. Cobelli, "Quantitative estimation of insulin sensitivity," *Am J Physiol*, vol. 236, pp. E667-77, Jun 1979.
- [50] C. Dalla Man, R. A. Rizza, and C. Cobelli, "Meal simulation model of the glucose-insulin system," *IEEE Trans Biomed Eng*, vol. 54, pp. 1740-9, Oct 2007.
- [51] R. Hovorka, V. Canonico, L. J. Chassin, U. Haueter, M. Massi-Benedetti, M. Orsini Federici, *et al.*, "Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes," *Physiol Meas*, vol. 25, pp. 905-20, Aug 2004.
- [52] A. R. Sedaghat, A. Sherman, and M. J. Quon, "A mathematical model of metabolic insulin signaling pathways," *Am J Physiol Endocrinol Metab*, vol. 283, pp. E1084-101, Nov 2002.
- [53] C. K. Ho, L. Rahib, J. C. Liao, G. Sriram, and K. M. Dipple, "Mathematical modeling of the insulin signal transduction pathway for prediction of insulin sensitivity from expression data," *Mol Genet Metab*, vol. 114, pp. 66-72, Jan 2015.
- [54] K. Zarkogianni, K. Mitsis, E. Litsa, M. T. Arredondo, G. Fico, A. Fioravanti, *et al.*, "Comparative assessment of glucose prediction models for patients with type 1 diabetes mellitus applying sensors for glucose and physical activity monitoring," *Med Biol Eng Comput*, vol. 53, pp. 1333-43, Dec 2015.
- [55] G. Cappon, L. Meneghetti, F. Prendin, J. Pavan, G. Sparacino, S. Del Favero, *et al.*, *A Personalized and Interpretable Deep Learning Based Approach to Predict Blood Glucose Concentration in Type 1 Diabetes*, 2020.
- [56] C. Mosquera-Lopez and P. G. Jacobs, "Incorporating Glucose Variability into Glucose Forecasting Accuracy Assessment Using the New Glucose Variability Impact Index and the Prediction Consistency Index: An LSTM Case Example," *J Diabetes Sci Technol*, vol. 16, pp. 7-18, Jan 2022.
- [57] C. Toffanin, F. Iacono, and L. Magni, *Personalized LSTM-based alarm systems for hypoglycemia prevention*, 2023.
- [58] M. Jensen, C. Dethlefsen, P. Vestergaard, and O. Hejlesen, "Prediction of Nocturnal Hypoglycemia From Continuous Glucose Monitoring Data in People With Type 1 Diabetes: A Proof-of-Concept Study," *Journal of Diabetes Science and Technology*, vol. 14, p. 193229681986872, 08/08 2019.
- [59] C. Mosquera-Lopez, R. Dodier, L. Wilson, J. El Youssef, J. Castle, P. Jacobs, *et al.*, "Predicting and Preventing Nocturnal Hypoglycemia in Type 1 Diabetes Using Big Data Analytics and Decision Theoretic Analysis," *Diabetes Technology & Therapeutics*, vol. 22, 04/16 2020.

- [60] A. Parcerisas, I. Contreras, A. Delecourt, A. Bertachi, A. Beneyto, I. Conget, *et al.*, "A Machine Learning Approach to Minimize Nocturnal Hypoglycemic Events in Type 1 Diabetic Patients under Multiple Doses of Insulin," *Sensors*, vol. 22, p. 1665, 02/21 2022.
- [61] V. Berikov, O. Kutnenko, J. Semenova, and V. Klimontov, "Machine Learning Models for Nocturnal Hypoglycemia Prediction in Hospitalized Patients with Type 1 Diabetes," *Journal of Personalized Medicine*, vol. 12, p. 1262, 07/31 2022.
- [62] C. Marling and R. Bunescu, "The OhioT1DM Dataset for Blood Glucose Level Prediction: Update 2020," *CEUR Workshop Proc*, vol. 2675, pp. 71-74, Sep 2020.
- [63] A. Gelman, F. Bois, and J. Jiang, "Physiological Pharmacokinetic Analysis Using Population Modeling and Informative Prior Distributions," *Journal of the American Statistical Association*, vol. 91, 04/11 1997.
- [64] C. Dalla Man, M. Camilleri, and C. Cobelli, "A System Model of Oral Glucose Absorption: Validation on Gold Standard Data," *IEEE transactions on bio-medical engineering*, vol. 53, pp. 2472-8, 01/01 2007.
- [65] C. Dalla Man, D. M. Raimondo, R. A. Rizza, and C. Cobelli, "GIM, simulation software of meal glucose-insulin model," *J Diabetes Sci Technol*, vol. 1, pp. 323-30, May 2007.
- [66] M. Sjöstrand, A. Holmång, and P. Lönnroth, "Measurement of interstitial insulin in human muscle," *Am J Physiol*, vol. 276, pp. E151-4, Jan 1999.
- [67] C. Raffel and D. Ellis, "Feed-Forward Networks with Attention Can Solve Some Long-Term Memory Problems," 12/29 2015.
- [68] o. pramod. (2023). *Cross Validation*. Available: <https://medium.com/@ompramod9921/cross-validation-623620ff84c2>
- [69] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning," *Proceedings of The 33rd International Conference on Machine Learning*, 06/06 2015.
- [70] S. Mougiakakou, C. Bartsocas, E. Bozas, N. Chaniotakis, D. Iliopoulou, I. Kouris, *et al.*, "SMARTDIAB: A Communication and Information Technology Approach for the Intelligent Monitoring, Management and Follow-up of Type 1 Diabetes Patients," *IEEE transactions on information technology in biomedicine : a publication of the IEEE Engineering in Medicine and Biology Society*, vol. 14, pp. 622-33, 05/01 2010.
- [71] M. Rybicka, R. Krysiak, and B. Okopien, "The dawn phenomenon and the Somogyi effect - Two phenomena of morning hyperglycaemia," *Endokrynologia Polska*, vol. 62, pp. 276-84, 05/01 2011.
- [72] A. Güemes González, G. Cappon, B. Hernandez, M. Reddy, N. Oliver, P. Georgiou, *et al.*, "Predicting Quality of Overnight Glycaemic Control in Type 1 Diabetes Using Binary Classifiers," *IEEE Journal of Biomedical and Health Informatics*, vol. PP, pp. 1-1, 09/13 2019.
- [73] I. Afentakis, R. Unsworth, P. Herrero, N. Oliver, M. Reddy, and P. Georgiou, *Prediction of Nocturnal Hypoglycaemia in Adults with Type 1 Diabetes using Machine Learning Classifiers*, 2023.
- [74] B. Naderalvojud and T. Hernandez-Boussard, "Improving machine learning with ensemble learning on observational healthcare data," *AMIA ... Annual Symposium proceedings. AMIA Symposium*, vol. 2023, pp. 521-529, 01/11 2024.
- [75] M. Athanasiou, "Development of interpretable machine learning models to support diabetes management," Doctoral, National Technical University of Athens, 2023.
- [76] K. Gasmi, L. Ben Ammar, H. Elshammari, and F. Yahya, "Prediction of Uncertainty Estimation and Confidence Calibration Using Fully Convolutional Neural Network," *Computers, Materials & Continua*, vol. 75, pp. 2557-2573, 03/27 2023.
- [77] A. Salih, Z. Raisi, I. Boscolo Galazzo, P. Radeva, S. Petersen, G. Menegaz, *et al.*, *Commentary on explainable artificial intelligence methods: SHAP and LIME*, 2023.

- [78] N. M. M. Athanasiou, K. Zarkogianni, K. S. Nikita, "Development of Personalized Interpretable Multilevel Prediction Models for the Risk Assessment of Hypoglycemia in Type 1 Diabetes," presented at the EMBC, 2024.