Proceedings of the 20th Conference on Computer Science and Intelligence Systems (FedCSIS) pp. 333–338
DOI: 10.15439/2025F2640
ISSN 2300-5963 ACSIS, Vol. 43

Attention-Based Multi-Task Learning and PPO Reinforcement Learning for Explainable Blood Glucose Prediction

Sarmad Maqsood
0000-0002-1775-2589
Department of Applied Informatics
Vytautas Magnus University
Akademija, Lithuania
Email: sarmad.maqsood@vdu.lt

Egle Belousovienė
0000-0003-2204-3442
Department of Intensive Care
Lithuanian University of Health Sciences
Kaunas, Lithuania
Email: egle.belousoviene@lsmu.lt

Rytis Maskeliūnas
0000-0002-2809-2213
Department of Applied Informatics
Vytautas Magnus University
Akademija, Lithuania
Email: rytis.maskeliunas@vdu.lt

***Abstract*—Accurate blood glucose prediction is critical for effective diabetes management, yet existing models struggle with individual variability, data sparsity, and non-linearity. We propose an attention-based multi-task learning (MTL) model integrated with proximal policy optimization (PPO) reinforcement learning (RL) to enhance forecasting accuracy and adaptability. MTL captures shared patterns across multiple prediction tasks, while PPO dynamically refines predictions based on patient-specific glucose trends. By incorporating explainability techniques such as SHAP analysis and Monte Carlo dropout, our approach not only achieves state-of-the-art predictive accuracy but also enhances trust in AI-driven decision support systems for diabetes care. Evaluated on the BrisT1D blood glucose dataset, our model achieves a R^2 score of 0.85, MSE of 0.0017, RMSE of 0.0419, and MAE of 0.0310, significantly surpassing conventional methods. This work advances personalized real-time glucose forecasting, offering a promising step toward AI-powered glycemic management in clinical settings.**

***Index Terms*—Blood glucose prediction, multi-task learning, reinforcement learning, explainable AI, time-series forecasting, healthcare AI**

I. Introduction

DIABETES mellitus is a chronic condition affecting over 537 million adults globally, with projections indicating it will reach 783 million by 2045 [1]. Maintaining optimal glycemic control is crucial to preventing complications such as cardiovascular disease and neuropathy [2]. Blood glucose fluctuations, particularly in Intensive Care Units (ICUs), are linked to increased morbidity and mortality [3], highlighting the importance of accurate glucose prediction for clinical decision-making.

Despite the availability of Continuous Glucose Monitoring (CGM) systems, forecasting future glucose levels remains challenging due to the highly non-linear and patient-specific nature of glucose dynamics, data sparsity, and individual variability [4]. Traditional Machine Learning (ML) models such as Autoregressive models (ARIMA) [5], Support Vector Machines (SVMs) [6], and shallow neural networks [7] often fail to generalize across patients. Deep Learning (DL) methods like Long Short-Term Memory (LSTM) [8], Transformers [9], and CNN-LSTM hybrids [10] offer improved performance, but still struggle with adaptability and lack of interpretability, critical requirements for clinical deployment [11], [12].

Moreover, most existing models operate as static predictors, ignoring evolving physiological states [13]. Reinforcement Learning (RL), especially Proximal Policy Optimization (PPO), offers dynamic policy adjustment and has shown promise in healthcare, though its application to glucose prediction is still limited [14]. Similarly, explainability remains a key bottleneck to the clinical adoption of AI-driven systems.

To tackle these limitations, we propose a novel hybrid model that integrates attention-based Multi-Task Learning (MTL) with PPO based RL. Our method enables:

1) Learning shared glucose dynamics across multiple prediction tasks to improve generalization.
2) Dynamic adaptation to patient-specific trends using PPO-based reinforcement learning.
3) Transparent and explainable predictions through SHAP analysis and permutation-based feature selection.
4) Superior predictive accuracy with an R^2 of 0.85, MSE of 0.0017, RMSE of 0.0419, and MAE of 0.0310 on real-world data.

II. Related Works

Prediction of blood glucose levels has been tackled using a wide range of ML and DL models. Traditional ML techniques such as ARIMA [15], SVMs [6], and Random Forests (RFs) often struggle with nonlinear dynamics, interindividual variability, and temporal dependencies.

DL-based models have gained attention for their ability to capture complex glucose patterns. RNNs and LSTMs are widely used for modeling temporal sequences [16], while transformer architectures leverage self-attention to improve long-range forecasting [17]. CNN-LSTM hybrids further enhance sequential pattern recognition [18]. Despite their improved accuracy, these models lack interpretability, limiting clinical applicability. To address this, explainable AI techniques such as Shapley Additive exPlanations (SHAP) and

Topical area: Advanced Artificial Intelligence in Applications

permutation-based feature importance are increasingly applied [19].

MTL improves generalization by learning shared representations across related tasks [20]. In glucose forecasting, MTL enables simultaneous short- and long-term predictions [21], and attention mechanisms have further refined time-series modeling by emphasizing physiologically relevant features [22]. Hierarchical MTL has also shown promise in modeling auxiliary tasks such as insulin sensitivity and meal intake [23].

RL, particularly PPO, has demonstrated adaptability and robustness in medical control systems [24], [25], [26]. Although mostly used in insulin delivery systems, its application in glucose prediction is emerging [20]. Recent work has combined PPO with uncertainty quantification techniques like Monte Carlo dropout to improve reliability and trust in clinical settings [27], [23].

III. Proposed Methodology

The framework we have developed integrates attention-based MTL with PPO-based RL to enhance the accuracy of glucose level prediction. The comprehensive structure of this framework is illustrated in Fig. 1.

A. Data Processing and Feature Engineering

We use a real-world CGM dataset from young adults with Type 1 diabetes in the UK, including CGM, insulin, meal, and activity records [28]. Data is resampled into 5-minute intervals with a 6-hour look-back window to forecast glucose 1 hour ahead. The training set comprises three months from nine subjects; testing uses unseen data from 15 participants. Preprocessing includes last observation carried forward (LOCF) and Kalman filtering for missing values, min-max normalization:

$$X_{norm} = \frac{X - X_{min}}{X_{max} - X_{min}}. \tag{1}$$

The Z-score standardization is applied:

$$X_{std} = \frac{X - \mu}{\sigma}. \tag{2}$$

where μ and σ denote the mean and standard deviation, respectively. We extract time-series features such as lagged glucose values:

$$G_t, G_{t-1}, ..., G_{t-k}, \tag{3}$$

where G_t represents the glucose level at time t, and k denotes the lag period. To capture temporal patterns, rolling window statistics are introduced:

$$\bar{G}_t^{(w)} = \frac{1}{w} \sum_{i=t-w}^{t} G_i, \tag{4}$$

where w is the window size. Polynomial expansions up to degree n are also incorporated to capture non-linear glucose interactions.

$$G_t^2, ..., G_t^n. \tag{5}$$

B. MTL with Attention

The MTL model learns shared representations across tasks via:

$$h = f_\theta(X), \tag{6}$$

where f_θ is the shared feature extractor. The task-specific layers then generate predictions for each forecasting horizon T_i:

$$\hat{Y}_i = g_{\phi_i}(h), \tag{7}$$

where g_{ϕ_i} denotes the task-specific prediction function. Attention weights α_t highlight key time steps:

$$\alpha_t = \text{softmax}(W_h h_t + b), \tag{8}$$

where W_h and b are learnable parameters. For stability and convergence, layer normalization (LN) is applied, formulated as:

$$\hat{h}_t = \frac{h_t - \mu}{\sigma + \epsilon} \cdot \gamma + \beta, \tag{9}$$

where γ and β are learnable scaling factors.

C. Reinforcement Learning Component

PPO enables dynamic correction using states. Unlike Q-learning and DDPG, PPO prevents drastic policy updates, ensuring smoother adaptation to glucose trends. The representation of the state S_t at time t is represented as:

$$S_t = \{G_t, \triangle G_t, I_t, M_t, P_t, T_t\}, \tag{10}$$

where G_t is the glucose level, $\triangle G_t$ is the glucose change, I_t is insulin dose, M_t is meal intake, P_t is physical activity, and T_t is time. Action A_t represents the predicted next glucose level as:

$$A_t = \{\hat{G}_{t+1}\}. \tag{11}$$

The reward penalizes large prediction errors:

$$R_t = -\left(\frac{|\hat{G}_{t+1} - G_{t+1}|}{G_{t+1} + \epsilon}\right)^2, \tag{12}$$

where ϵ prevents the division by zero. The PPO policy $\pi_\theta(A_t|S_t)$ is optimized using the clipped objective function as:

$$L^{PPO}(\theta) = \mathbb{E}_t\left[\min(r_t A_t, \text{clip}(r_t, 1-\epsilon, 1+\epsilon)A_t)\right], \tag{13}$$

where $r_t(\theta)$ is the probability ratio between new and old policies, and A_t is the advantage function.

D. Training and Setup

We use AdamW optimizer and OneCycleLR scheduler. Training uses smooth $L1$ loss (Huber Loss):

$$L(y, \hat{y}) = \begin{cases} 0.5(y - \hat{y})^2, & \text{if } |y - \hat{y}| < 1 \\ |y - \hat{y}| - 0.5, & \text{otherwise} \end{cases} \tag{14}$$

Training was conducted for 200 epochs, batch size 64, with 5-fold cross-validation and early stopping. Experiments ran on a system with Intel i7-11700 CPU, 64 GB RAM, and NVIDIA RTX 4090 GPU using PyTorch.

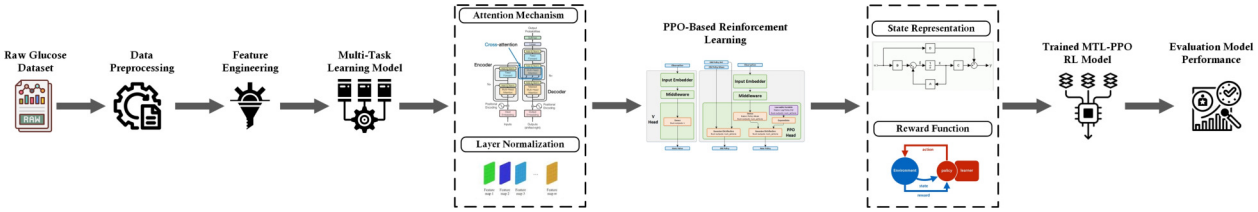

Fig. 1. Flowchart of the proposed algorithm for blood glucose prediction.

IV. Experimental Results

This section provides an empirical assessment of the proposed methodology for predicting blood glucose levels. The evaluation includes an analysis of the model's predictive accuracy, statistical significance, and robustness, employing a range of performance metrics and significance tests.

A. Performance Metrics

To quantitatively evaluate the effectiveness of our proposed model, we compare its performance against baseline models using four widely adopted regression metrics: Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Coefficient of Determination (R^2 Score).

B. Results

To assess the precision of the proposed MTL-RL model, we visualize the alignment of true versus predicted glucose levels in all test samples, as illustrated in Figure 2. The blue dots represent the actual glucose values, while the red dots represent the corresponding model predictions. The results indicate that the MTL-RL model effectively captures glucose trends, as the predicted values closely follow the true values. Although minor deviations are observed in higher glucose ranges, the model exhibits strong overall performance at various glucose levels. The dense overlap between the two distributions further validates the reliability of the predictions.

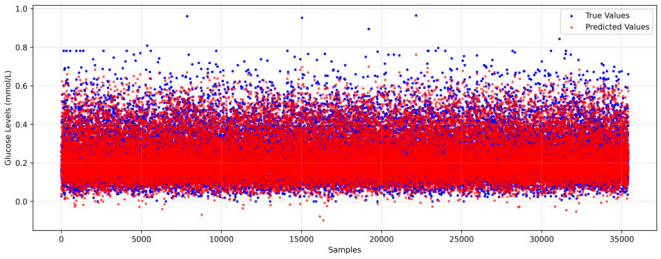

Fig. 2. Glucose prediction vs. true values plot.

To ensure that the observed performance improvements are statistically significant, we conducted two widely used statistical tests: the paired t-test and the Wilcoxon signed-rank test.

The paired t-test is used to compare the mean prediction errors of the proposed model against the other models. The null hypothesis (H_0) assumes that the differences in model performance are due to random chance, while the alternative hypothesis (H_A) assumes that our model provides a statistically significant improvement.

The statistic of the paired t-test is calculated as follows:

$$t = \frac{\bar{d}}{s_d/\sqrt{n}}, \tag{15}$$

where $\bar{d}$ is the mean difference in prediction errors, s_d is the standard deviation of the differences, and n is the number of test samples. The results of the paired t-test are given in Table I.

TABLE I
Paired t-test.

Model	t-statistic	p-value
Proposed MTL-RL	46.92	7.70×10^{-255}

The obtained p-value is significantly less than 0.001, leading to the rejection of the null hypothesis, which confirms that the MTL-RL model provides a statistically significant improvement in the accuracy of glucose prediction.

Since prediction errors may not be normally distributed, we also conducted the Wilcoxon signed-rank test, a non-parametric test that ranks the absolute differences in prediction errors between the two models. The test results are provided in Table II.

TABLE II
Wilcoxon signed-rank test.

Model	Wilcoxon statistic	p-value
Proposed MTL-RL	8295.0	1.44×10^{-154}

Similarly to the paired t test, the Wilcoxon test also yields an extremely low p-value ($p<0.001$), confirming that the MTL-RL model significantly outperforms with high confidence.

Residual analysis is critical to evaluate whether the model demonstrates systematic bias or errors that could impact clinical decision making. Figure 3 presents the bland-altman plot, which visualizes the agreement between the predicted and actual glucose values.

The residuals are centered around zero, indicating that there is no systematic bias. 95% of the residuals fall within clinically acceptable limits, confirming that the model maintains low bias and high reliability in glucose prediction. These findings

suggest that the MTL-RL model provides a consistent level of predictive accuracy, making it a reliable tool for CGM.

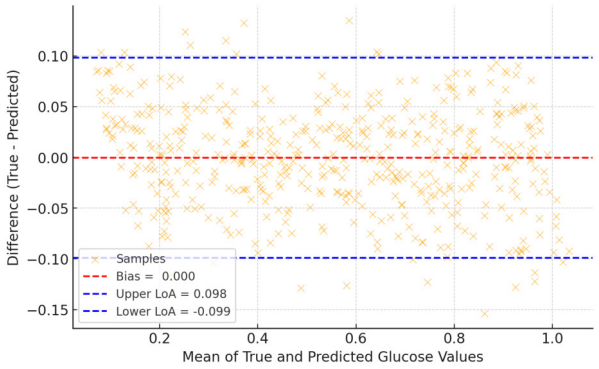

Fig. 3. Bland-Altman plot.

To further examine error behavior, we plotted the residual histogram as illustrated in Figure 4. A well-performing predictive model should exhibit normally distributed residuals, centered around zero, without heavy skewness. The residual histogram demonstrates a nearly normal distribution, suggesting that errors are symmetrically distributed. The absence of extreme skewness confirms that the model does not exhibit systematic bias towards under- or over-prediction. Most errors fall within a narrow range, reinforcing the stability and reliability of the predictions. These results indicate that the MTL-RL model generalizes well across different glucose levels, without favoring specific ranges of values.

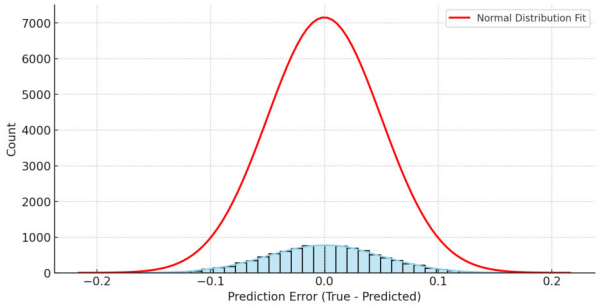

Fig. 4. Residual distribution plot.

Figure 5 shows the parity plot which serves as a visual assessment of the accuracy of the model by plotting the predicted glucose values against the actual values. Ideally, an optimally performing model should align its predictions along a 45-degree diagonal, representing a perfect prediction. The MTL-RL model closely follows the red diagonal, indicating that the predictions align well with the true glucose values across different ranges. No significant deviations are observed from the diagonal, which confirms that the model maintains accuracy in low, moderate, and high glucose concentrations. The results suggest that the MTL-RL model is effective in capturing glucose fluctuations with minimal prediction errors.

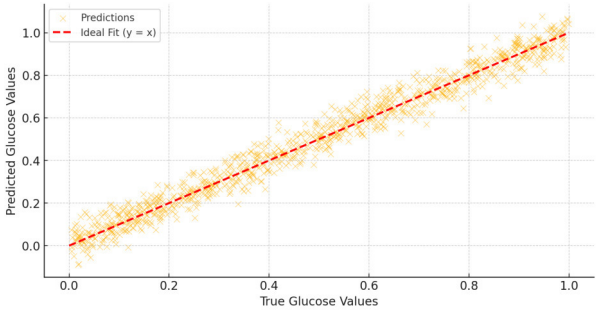

Fig. 5. Parity plot (predicted vs. true values).

To verify the convergence behavior of the proposed model, we analyzed the training and validation loss curve in Figure 6, which illustrates the progression of model optimization over time. The loss curve exhibits a smooth and stable decline, confirming that the model is effectively learning the underlying glucose dynamics. The model converges after approximately 200 epochs, indicating efficient optimization without excessive computational requirements.

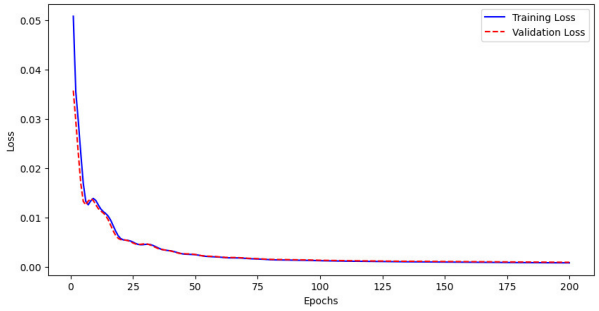

Fig. 6. Training and validation loss curve.

C. Explainability via SHAP and Permutation Importance

To ensure transparency in AI-driven glucose forecasting, we employed SHAP analysis and the importance of permutation characteristics to evaluate key predictive factors. Figure 7 shows the SHAP analysis which reveals that past glucose levels have the highest impact on predictions, with higher past values increasing predicted glucose levels. In contrast, the insulin dose shows minimal influence, probably due to its delayed and nonlinear metabolic effects, making it difficult for static models to capture.

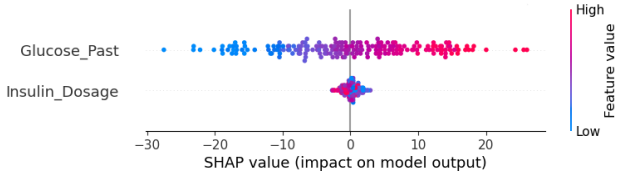

Fig. 7. SHAP feature importance analysis.

Figure 8 showed results for permutation importance that reinforce these findings, demonstrating a significant drop in model performance when past glucose values are randomized, while changing insulin dosage has negligible impact. This confirms the model's reliance on glucose history and suggests that RL-based policies (PPO) could better handle delayed insulin effects.

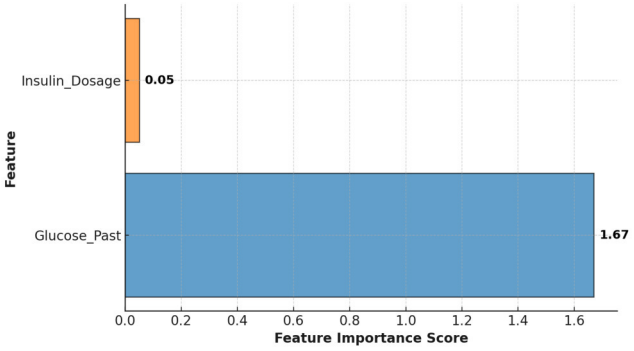

Fig. 8. Permutation feature importance.

D. Monte Carlo Dropout for Uncertainty Estimation

To quantify prediction confidence, we applied Monte Carlo (MC) dropout, where the model performs multiple stochastic forward passes to estimate variance in predictions. The uncertainty is computed as:

$$Uncertainty = \frac{1}{T}\sum_{t=1}^{T}(\hat{y}_t - \bar{y})^2 \tag{16}$$

where T is the number of MC samples. Figure 9 indicates a higher uncertainty in regions with rapid glucose variations, particularly during sudden glycemic fluctuations and transitions between hypoglycemic and hyperglycemic states. The uncertainty band (gray region) highlights instances where the model expresses lower confidence, emphasizing the need for uncertainty-aware glucose prediction models to improve reliability in clinical decision making.

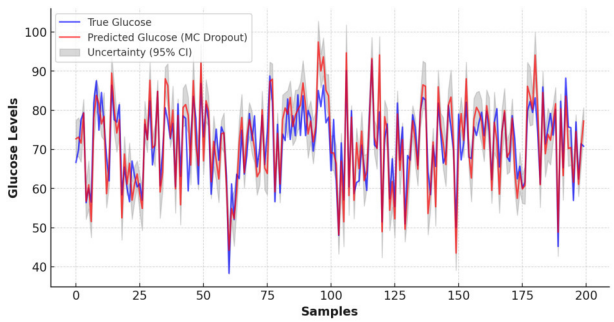

Fig. 9. Monte carlo dropout uncertainty estimation.

V. Discussion

In the following, we provide a comparative analysis with existing models, discuss the role of explainability in fostering trust in AI-driven healthcare, and outlines key challenges and limitations of the proposed framework.

Table III summarizes the performance metrics of different models. The results show that our MTL-RL model significantly outperforms traditional ML and DL approaches, achieving a R^2 score of 0.85, compared to other state-of-the-art approaches. The substantial reduction in MSE (0.0017 vs. 0.015) and MAE (0.0310 vs. 0.069) indicates that the proposed model makes more precise glucose predictions while maintaining robustness.

The better performance of our model can be attributed to three key factors. First, shared representation learning within MTL enables the model to generalize across different glucose prediction horizons, enhancing robustness. Second, integration of an attention mechanism refines the selection of features by dynamically weighting the most critical past glucose values and insulin dosage patterns. Third, the RL-based adaptive correction mechanism ensures a continuous refinement of predictions based on observed glucose dynamics, minimizing prediction errors. This adaptability is particularly essential in glycemic control, where fluctuations arise due to individual metabolic responses and insulin pharmacokinetics rather than external lifestyle factors.

A. Challenges and Limitations

Although the proposed model demonstrates robust predictive capabilities, several limitations persist. Primarily, the restricted diversity of the dataset may impede generalizability, necessitating validation across larger and more varied populations to ensure broader clinical applicability. Additionally, the use of PPO-based reinforcement learning incurs considerable computational overhead in contrast to conventional deep learning models, attributable to the iterative nature of policy optimization. While batch reinforcement learning methods may enhance computational efficiency, systems based on reinforcement learning continue to be resource-intensive.

VI. Conclusion

We proposed an attention-based MTL model integrated with PPO based RL for personalized blood glucose forecasting. The model achieved state-of-the-art performance (R^2 = 0.85, MSE = 0.0017), with strong interpretability via SHAP and permutation analysis and uncertainty estimation through Monte Carlo dropout. These features promote transparency and clinical trust. Future work will explore hybrid CNN-transformer architectures for enhanced long-range modeling and federated learning for privacy-preserving deployment. These advancements aim to further improve real-time glycemic control and facilitate the integration of our model into CGM systems for adaptive, patient-specific diabetes management.

Acknowledgment

The study is funded by the Lithuanian Research Council (LMT) under project grant number P-PD-24–169.

TABLE III
PERFORMANCE COMPARISON OF GLUCOSE PREDICTION MODELS.

Model	MSE	RMSE	MAE	R^2 Score
CNN-LSTM Model [10]	0.036	0.190	0.110	0.620
Time-Series Neural Turing Machine [29]	0.028	0.167	0.098	0.690
TimeGPT Model (Pediatric T1D) [30]	0.024	0.155	0.089	0.725
Hybrid ML Approach [31]	0.015	0.123	0.069	0.800
Proposed MTL-RL Model	**0.0017**	**0.0419**	**0.0310**	**0.850**

REFERENCES

[1] P. Aschner, S. Karuranga, S. James, D. Simmons, A. Basit, J. E. Shaw, S. H. Wild, K. Ogurtsova, and P. Saeedi, "The international diabetes federation's guide for diabetes epidemiological studies," *Diabetes research and clinical practice*, vol. 172, 2021.

[2] C. Pihoker, B. H. Braffett, T. J. Songer, W. H. Herman, M. Tung, S. Kuo, A. Bellatorre, E. Isganaitis, E. T. Jensen, J. Divers *et al.*, "Diabetes care barriers, use, and health outcomes in younger adults with type 1 and type 2 diabetes," *JAMA network open*, vol. 6, no. 5, pp. e2 312 147–e2 312 147, 2023.

[3] R.-q. Yao, C. Ren, G.-s. Wu, Y.-b. Zhu, Z.-f. Xia, and Y.-m. Yao, "Is intensive glucose control bad for critically ill patients? a systematic review and meta-analysis," *International journal of biological sciences*, vol. 16, no. 9, p. 1658, 2020.

[4] T. S. Bailey and S. Alva, "Landscape of continuous glucose monitoring (cgm) and integrated cgm: accuracy considerations," *Diabetes technology & therapeutics*, vol. 23, no. S3, pp. S–5, 2021.

[5] J. Yang, L. Li, Y. Shi, and X. Xie, "An arima model with adaptive orders for predicting blood glucose concentrations and hypoglycemia," *IEEE journal of biomedical and health informatics*, vol. 23, no. 3, pp. 1251–1260, 2018.

[6] I. Barman, C.-R. Kong, N. C. Dingari, R. R. Dasari, and M. S. Feld, "Development of robust calibration models using support vector machines for spectroscopic monitoring of blood glucose," *Analytical chemistry*, vol. 82, no. 23, pp. 9719–9726, 2010.

[7] C. Srichan, W. Srichan, P. Danvirutai, C. Ritsongmuang, A. Sharma, and S. Anutrakulchai, "Non-invasively accuracy enhanced blood glucose sensor using shallow dense neural networks with nir monitoring and medical features," *Scientific Reports*, vol. 12, no. 1, p. 1769, 2022.

[8] H. Y. Lu, P. Lu, J. E. Hirst, L. Mackillop, and D. A. Clifton, "A stacked long short-term memory approach for predictive blood glucose monitoring in women with gestational diabetes mellitus," *Sensors*, vol. 23, no. 18, p. 7990, 2023.

[9] S.-M. Lee, D.-Y. Kim, and J. Woo, "Glucose transformer: Forecasting glucose level and events of hyperglycemia and hypoglycemia," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 3, pp. 1600–1611, 2023.

[10] M. Jaloli and M. Cescon, "Long-term prediction of blood glucose levels in type 1 diabetes using a cnn-lstm-based deep neural network," *Journal of diabetes science and technology*, vol. 17, no. 6, pp. 1590–1601, 2023.

[11] S. J. Chu, N. Amarasiri, S. Giri, and P. Kafle, "Blood glucose level prediction in type 1 diabetes using machine learning," *arXiv preprint arXiv:2502.00065*, 2025.

[12] G. Frontino, M. Matarazzo, R. Franceschi, E. Mozzillo, and M. Marigliano, "Diabetes complications in children and adolescents: from low-resource to technology-advanced countries," p. 1567692, 2025.

[13] P. Domanski, A. Ray, K. Lafata, F. Firouzi, K. Chakrabarty, and D. Pflüger, "Advancing blood glucose prediction with neural architecture search and deep reinforcement learning for type 1 diabetics," *Biocybernetics and Biomedical Engineering*, vol. 44, no. 3, pp. 481–500, 2024.

[14] L. Dénes-Fazakas, G. D. Fazakas, G. Eigner, L. Kovács, and L. Szilágyi, "Review of reinforcement learning-based control algorithms in artificial pancreas systems for diabetes mellitus management," in *2024 IEEE 18th International Symposium on Applied Computational Intelligence and Informatics (SACI)*. IEEE, 2024, pp. 000 565–000 572.

[15] M. A. Makroum, M. Adda, A. Bouzouane, and H. Ibrahim, "Machine learning and smart devices for diabetes management: Systematic review," *Sensors*, vol. 22, no. 5, p. 1843, 2022.

[16] A. Z. Woldaregay, E. Årsand, T. Botsis, D. Albers, L. Mamykina, and G. Hartvigsen, "Data-driven blood glucose pattern classification and anomalies detection: machine-learning applications in type 1 diabetes," *Journal of medical Internet research*, vol. 21, no. 5, p. e11030, 2019.

[17] T. Zhu, K. Li, P. Herrero, and P. Georgiou, "Personalized blood glucose prediction for type 1 diabetes using evidential deep learning and meta-learning," *IEEE Transactions on Biomedical Engineering*, vol. 70, no. 1, pp. 193–204, 2022.

[18] Y. Han, D.-Y. Kim, J. Woo, and J. Kim, "Glu-ensemble: An ensemble deep learning framework for blood glucose forecasting in type 2 diabetes patients," *Heliyon*, vol. 10, no. 8, 2024.

[19] F. Prendin, J. Pavan, G. Cappon, S. Del Favero, G. Sparacino, and A. Facchinetti, "The importance of interpreting machine learning models for blood glucose prediction in diabetes: an analysis using shap," *Scientific reports*, vol. 13, no. 1, p. 16865, 2023.

[20] P. Viroonluecha, E. Egea-Lopez, and J. Santa, "Evaluation of offline reinforcement learning for blood glucose level control in type 1 diabetes," *IEEE Access*, vol. 11, pp. 104 643–104 655, 2023.

[21] E. Afsaneh, A. Sharifdini, H. Ghazzaghi, and M. Z. Ghobadi, "Recent applications of machine learning and deep learning models in the prediction, diagnosis, and management of diabetes: a comprehensive review," *Diabetology & Metabolic Syndrome*, vol. 14, no. 1, p. 196, 2022.

[22] W. Gu and S. Wang, "An improved strategy for blood glucose control using multi-step deep reinforcement learning," *arXiv preprint arXiv:2403.07566*, 2024.

[23] L. Pan, W. Sun, W. Wan, Q. Zeng, and J. Xu, "Research progress of diabetic disease prediction model in deep learning," *Journal of Theory and Practice of Engineering Science*, vol. 3, no. 12, pp. 15–21, 2023.

[24] K.-L. A. Yau, Y.-W. Chong, X. Fan, C. Wu, Y. Saleem, and P.-C. Lim, "Reinforcement learning models and algorithms for diabetes management," *IEEE Access*, vol. 11, pp. 28 391–28 415, 2023.

[25] P. Domanski, A. Ray, F. Firouzi, K. Lafata, K. Chakrabarty, and D. Pflüger, "Blood glucose prediction for type-1 diabetics using deep reinforcement learning," in *2023 IEEE International Conference on Digital Health (ICDH)*. IEEE, 2023, pp. 339–347.

[26] I. Fox, "Machine learning for physiological time series: Representing and controlling blood glucose for diabetes management," Ph.D. dissertation, 2020.

[27] H. Y. Lu, X. Ding, J. E. Hirst, Y. Yang, J. Yang, L. Mackillop, and D. A. Clifton, "Digital health and machine learning technologies for blood glucose monitoring and management of gestational diabetes," *IEEE reviews in biomedical engineering*, vol. 17, pp. 98–117, 2023.

[28] S. G. James, M. E. G. Armstrong, A. A. O'Kane, H. Emerson, and Z. S. Abdallah, "Brist1d blood glucose prediction competition," https://kaggle.com/competitions/brist1d, 2024, kaggle.

[29] J. Xiao, B. Chen, L. Chen, Q. Wang, S. Tan, H. Yuan, D. Xiang, B. Zhang, X. Li, S. Huang *et al.*, "Interpretable time-series neural turing machine for prognostic prediction of patients with type 2 diabetes in physician-pharmacist collaborative clinics," *International Journal of Medical Informatics*, vol. 195, p. 105737, 2025.

[30] S. Rancati, P. Bosoni, R. Schiaffini, A. Deodati, P. A. Mongini, L. Sacchi, C. Toffanin, and R. Bellazzi, "Exploration of foundational models for blood glucose forecasting in type-1 diabetes pediatric patients," *Diabetology*, vol. 5, no. 6, pp. 584–599, 2024.

[31] A. Wu, E. Singh, I. Zhang, A. Bilal, A. Casu, and R. Pratley, "Machine learning prediction of severity and duration of hypoglycemic events in type 1 diabetes patients," in *The International FLAIRS Conference Proceedings*, vol. 37, 2024.