Neural Architecture Search for Blood Glucose Prediction in Type-1 Diabetics

Anthony Liardo*[1], Aritra Ray*[2], Farshad Firouzi[2], Kyle Lafata[2], and Krishnendu Chakrabarty[1]

[1]Arizona State University, *AZ, USA*
[2]Duke University, *NC, USA*
[1]{*aliardo, krishnendu.chakrabarty*}*@asu.edu*
[2]{*aritra.ray, farshad.firouzi, kyle.lafata*}*@duke.edu*

2024 IEEE 20th International Conference on Body Sensor Networks (BSN) | 979-8-3315-3014-3/24/$31.00 ©2024 IEEE | DOI: 10.1109/BSN63547.2024.10780668

***Abstract*—For subjects affected with type-1 diabetes mellitus, accurately predicting future blood glucose values helps regulate insulin delivery. This paper introduces a dual Q-network-based neural architecture search approach to develop and train personalized BG prediction models for individuals affected with type-1 diabetes mellitus. Utilizing historical blood glucose data collected via body sensor networks, the proposed model forecasts future blood glucose levels. When evaluated on the OhioT1DM dataset, the proposed approach shows significant improvements over the state-of-the-art, achieving a 46.78% reduction in root mean square error and a 56.05% reduction in mean absolute error while predicting blood glucose values 5 minutes into the future.**

***Index Terms*—Neural architecture search, Diabetes, Blood glucose**

I. Introduction

Type-1 Diabetes Mellitus (T1DM) is a pancreas-specific medical condition wherein the β cells lose the ability of insulin production, which is essential for regulating Blood Glucose (BG) in humans. Uncontrolled BG, in the absence of medical treatment, leads to complications such as diabetic retinopathy, and affects functionalities of other organs in the human body [1]. International Diabetes Federation (IDF) estimates that in 2022, about 8.75 million people were affected with T1DM, and is expected to rise in the future [2].

Given the historical BG values measured using body sensor-based glucose monitoring devices, this paper proposes a dual Q-network Neural Architecture Search (NAS) approach to design personalized BG prediction models for predicting future BG values in individuals affected by T1DM. The dual Q-network NAS model consists of a target Q-network and a parent Q-network. The parent Q-network randomly samples architectures for BG prediction model from the search space, or uses a prediction based on experience replay. The target Q-network is responsible for enabling a well-informed guided architecture selection for the BG prediction model based on the performance of the architectures sampled in prior. The model performance on the validation set is used as a reward signal to update the Q-network weights.

The paper evaluates the performance of the proposed dual Q-network NAS model over the OhioT1DM [3] dataset. The Continuous Glucose Monitoring (CGM) values for all 12 subjects in the dataset are recorded at 5-minute intervals. Based on historical BG values, this paper predicts future BG values at a Prediction Horizon (PH) of 5-minute. The proposed short-term BG prediction method demonstrates superior analytical performance, as evidenced by lower Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) in the PH, as well as superior clinical performance in regards to Surveillance Error Grid [4], in comparison to the state-of-the-art [5].

The rest of the paper is organized as follows: Section II focuses on the methodology, followed by the experimental setup in Section III. Section IV delves into the evaluation results, while Section V concludes the paper.

*Equal Contribution

II. Methodology

The proposed algorithm builds upon the framework of Chauhan et. al. [6] to sample neural networks from a defined search space, with further improvements. In traditional Q-networks [7], models often overestimate action rewards due to Q-value divergence. To address this, this paper employs a dual Q-network architecture consisting of a parent controller Q-network and a separate target Q-network. The target Q-network is responsible for enabling well-informed guided architecture selection, based on the performance of the architectures sampled in prior. Infrequently updating the target controller Q-network stabilizes the training process by maintaining stable Q-network weights for n consecutive epochs, wherein the target Q-network is updated every n epochs. On the other hand, the parent controller Q-network randomly samples architectures from the search space, while updating its network weights based on the reward signal in each epoch, or it uses its network weights to predict the next action (the next personalized BG prediction model architecture). This choice of architecture generation is based on the ϵ-greedy strategy. Every generated architecture, either randomly sampled via the parent Q-network, or sampled in a well-informed fashion from the search space, is trained on the BG prediction dataset. The model performance on the validation set is used as a reward signal to update the Q-networks.

Within the NAS framework, the parent Q-network randomly samples architectures from the search space, illustrating an exploration strategy. On the other hand, the target Q-network

is responsible for a well-informed guided architecture selection from the search space, based on the performance of the architectures sampled in prior, elucidating an exploitation strategy. To strike a balance between the exploration and exploitation, this paper uses an ϵ-greedy strategy. Through an exponential decay of ϵ, selection of architectures from a well-trained parent Q-network is gradually prioritized over random selection of architectures by the parent Q-network, once the target Q-network's stability enables the parent Q-network to effectively estimate state-action pairs. This dual Q-network strategy effectively acts as a single agent, optimizing the search to generate neural networks with a superior performance.

III. Experimental setup

A. Dataset

This study utilizes the OhioT1DM dataset [3], a well-studied resource in the field of BG value prediction research. This dataset comprises information from 12 individuals diagnosed with T1DM. Over eight weeks, the dataset captures CGM data, insulin doses (both bolus and basal), heart rate, galvanic skin response, skin temperature, air temperature, step count, self-reported meal times with carbohydrate estimates, and self-reported occurrences of exercise, sleep, work, stress, and illness for each of the 12 participants. Throughout the data collection phase, participants used either Medtronic 530G or 630G insulin pumps and utilized Medtronic Enlite CGM sensors. The dataset ensures patient anonymity by employing randomized identifiers. It encompasses 19 distinct features, with a particular emphasis on CGM blood glucose levels recorded at five-minute intervals. Notably, the dataset includes data from seven male and five female participants, spanning ages from 20 to 80 years. To remain consistent with the state-of-the-art [5], this work considers only historical BG values to predict future BG values.

B. Preprocessing

The pre-processing step involves z-score normalization for each subject in the OhioT1DM dataset [3] for the BG values. To prevent data leakage, the z-score normalization score obtained for the BG values in the train set, was used over the test set, in a patient-specific manner. Equation 1 highlights the z-score normalization, wherein distribution mean μ_x and standard deviation σ_x for BG values in the train set, are at a patient-specific level.

$$x' = \frac{x - \mu_x}{\sigma_x} \tag{1}$$

Clinically, disparities in HbA1C and mean BG levels have been observed across various racial groups, with notable variations within racial groups [8]. This normalization approach emphasizes the detection of potentially concerning fluctuations in BG values, while overlooking statistical elements in the data that lack meaningful relevance for CGM prediction.

C. Neural Architecture

This paper proposes personalized BG prediction models for every subject of the OhioT1DM [3] dataset. At every epoch, the architecture randomly sampled by the parent Q-network from the search space, or a well-informed guided architecture selection from the search space by the target Q-network, is trained for an individual subject of the OhioT1DM dataset. The validation error is used as a reward signal to update the parent Q-network weights and target Q-network weights at every epoch and at every 3 epochs, respectively. The batch size was set at one-sixth of the training set size. The total number of layers in the BG prediction model was set as 5, with 32 LSTM-layer units. The training for personalized BG prediction networks was set to 370 epochs with an early stopping callback to prevent over-fitting on the training set. Table I illustrates the search space. Table II imposes constraints to determine the validity of the randomly sampled architecture from the search space by the parent Q-network, or the well-informed guided architecture selection from the search space by the target Q-network. The implementations are publicly available[1].

TABLE I: Neural Architecture Search Space.

Parameter	Values
Layer Types	{convolutional, LSTM, pooling, flatten, dense}
LSTM-layer Activation Function	{sigmoid, hyperbolic tangent, ReLU, ELU, SELU, swish}
LSTM-layer Dropout	{0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9}
Convolutional Filters	{16, 32, 64, 96, 128, 160, 192, 224, 256}
Convolutional Filter Size	{2, 3, 4, 5}
Pooling Size	{2, 3, 4, 5}
Convolutional Stride	{1, 2, 3}
Pooling Stride	{1, 2, 3, 4, 5}
Learning Rate	{0.05, 0.07, 0.09, 0.105, 0.12, 0.135}

TABLE II: Constraints for Valid BG Prediction Model Architecture.

Parameter	Values
First Layer	{convolutional, separable convolutional, depthwise convolutional, convolutional transpose}
Second and Third Layer	{convolutional, separable convolutional, depthwise convolutional, convolutional transpose, max pooling, average pooling, global max pooling, global average pooling, LSTM}
Two Fully Connected Output Layers	{flatten followed by dense}

D. Evaluation metrics

Analytical accuracy: The performance of the BG prediction model generated by the proposed algorithm is evaluated analytically through RMSE and MAE metrics. Over the entire test set (t), the model predicts BG value, $\hat{y}$ at every 5-min PH, denoted by i. RMSE and MAE is calculated at every such PH in comparison to the ground truth, denoted by y_i, in Equation (2) and (3) respectively.

[1]https://github.com/AnthonyLia/dqnas-for-ohiot1dm

$$\text{RMSE} = \sqrt{\frac{1}{t}\sum_{i=1}^{t}(y_i - \hat{y}_i)^2} \tag{2}$$

$$\text{MAE} = \frac{1}{t}\sum_{i=1}^{t}|y_i - \hat{y}_i| \tag{3}$$

Clinical accuracy: The clinical risk of the BG prediction model generated by the proposed algorithm is evaluated via the Surveillance Error Grid [4]. The error grid displays risk zones with clinical impact scores ranging from 0 (none) to 4 (extreme). Fig. 1 illustrates a simplified SEG with limits from 0 to 600 mg/dl and risk zones in 120 mg/dl intervals. The colors indicate the average risk rating, reflecting the consensus of expert opinions [4].

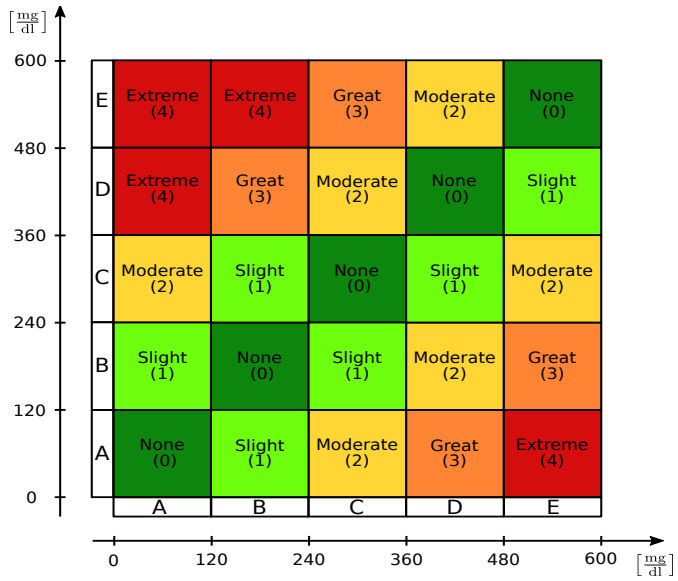

Fig. 1: Visualization of SEG risk zones.

IV. Evaluation Results

Single-Step Prediction for 5 minutes: The proposed Deep RL with NAS method predicts BG levels for one step in the future (corresponding to the next 5 minutes). Individual prediction models are trained and tested for all 12 patients in the OhioT1DM dataset [3]. The model's analytical performance is assessed using RMSE and MAE. As highlighted in Table III, the average RMSE and MAE across all 12 subjects are 5.266 and 3.331 respectively. In comparison to the state-of-the-art [5], the model achieves an average improvement in RMSE and MAE of 46.782% and 56.049%, respectively, for all 12 subjects of the OhioT1DM dataset [3]. With regards to RMSE and MAE, the proposed method has the best performance (corresponding to the lowest RMSE and MAE among the 12 subjects) on Patient ID 552, the worst performance (corresponding to the highest RMSE and MAE among the 12 subjects) on Patient ID 584, as inferred from Table III. Considering the two extreme cases, we provide visualizations of the predicted BG values against the ground truth on the test data for patient ID 552 (best-case) and patient ID 584 (worst-case), in Fig. 2.

Surveillance Error Grid for Single-Step Prediction: The clinical accuracy of the proposed method is assessed using the SEG. The SEG values for all the 12 subjects of the OhioT1DM [3] dataset is presented in Table IV. On average for all the 12 subjects of the OhioT1DM [3] dataset, 97.654%, 2.323%, 0.023%, 0.000%, and 0.000% of the predicted BG values fall within the no risk, slight risk, moderate risk, great risk, and extreme risk zone respectively when compared to the ground truth. Negligible predictions in the moderate risk zone, and no predictions in the great or extreme risk illustrate the efficacy of the proposed method. Due to the absence of a baseline for SEG in the 5-minute prediction horizon, the authors refrain from presenting any baseline comparisons. Fig. 3 shows our predictions overlaid on a continuously color-coded error grid.

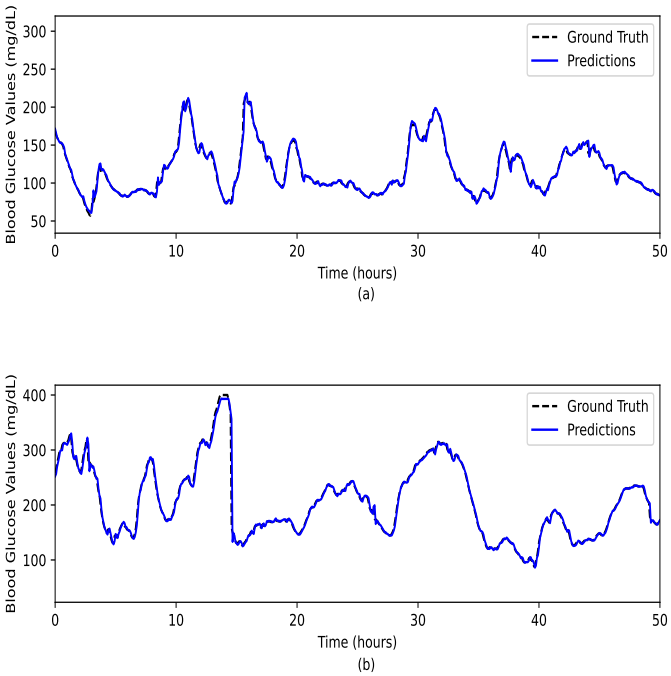

Fig. 2: Blood glucose ground truth in comparison to the prediction for a representative sample of the testing dataset of 50 hours for patient ID 552 (best-case) (a) and patient ID 584 (worst-case) (b)

V. Conclusion

This paper presents a dual Q-network based NAS to design and train personalized BG prediction models for subjects affected with T1DM. The models are evaluated analytically and clinically on the OhioT1DM dataset. For a prediction horizon of 5-minute, the predicted models generated through the proposed NAS approach demonstrates significant improvements over the state-of-the-art. In future work, the authors aim to validate the NAS strategy for longer prediction horizons of 30 and 45-minute, while also evaluating on additional T1DM datasets.

TABLE III: Single-Step Prediction Results (PH=5 mins)

Patient ID	Root Mean Square Error			Mean Absolute Error		
	Our Proposed Method	Baseline [5] Deep RL	% Improvement Over Baseline	Our Proposed Method	Baseline [5] Deep RL	% Improvement Over Baseline
540	4.923	12.270	59.877	3.383	9.441	64.167
544	4.021	8.806	54.337	3.044	7.087	57.048
552	3.435	8.336	58.793	2.962	6.573	54.936
559	6.114	10.920	44.011	3.264	8.293	60.642
563	5.620	9.017	37.673	3.281	6.907	52.497
567	7.166	9.309	23.021	3.599	7.293	50.651
570	4.066	10.940	62.834	2.913	8.626	66.230
575	5.804	11.640	50.137	3.038	8.914	65.919
584	7.769	11.630	33.199	3.876	8.745	55.677
588	4.202	8.542	50.807	3.504	6.894	49.173
591	5.170	9.840	47.459	3.905	7.225	45.952
596	4.903	8.070	39.244	3.199	6.360	49.701
Average	5.266	9.433	46.782	3.331	7.696	56.049

TABLE IV: SEG in Single-Step Prediction Results (PH=5 mins)

Patient ID	None (0)	Slight (1)	Moderate (2)	Great (3)	Extreme (4)
540	97.163	2.837	0.000	0.000	0.000
544	98.192	1.808	0.000	0.000	0.000
552	97.583	2.417	0.000	0.000	0.000
559	98.485	1.435	0.080	0.000	0.000
563	96.841	3.159	0.000	0.000	0.000
567	97.608	2.350	0.042	0.000	0.000
570	98.576	1.314	0.110	0.000	0.000
575	97.755	2.245	0.000	0.000	0.000
584	97.367	2.595	0.038	0.000	0.000
588	97.558	2.442	0.000	0.000	0.000
591	96.986	3.014	0.000	0.000	0.000
596	97.735	2.265	0.000	0.000	0.000
Average	97.654	2.323	0.023	0.000	0.000

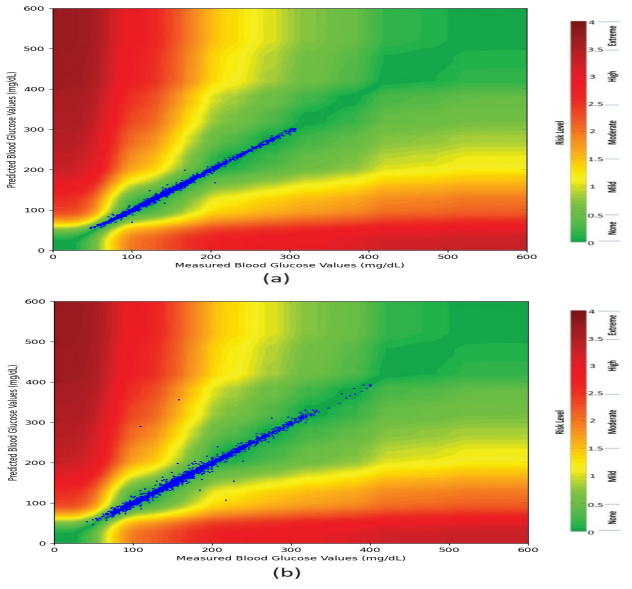

Fig. 3: SEG for Blood glucose predictions in comparison to the prediction for a representative sample of the testing dataset of 50 hours for patient ID 552 (best-case) (a) and patient ID 584 (worst-case) (b)

REFERENCES

[1] 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.

[2] H. Sun, P. Saeedi, S. Karuranga, M. Pinkepank, K. Ogurtsova, B. B. Duncan, C. Stein, A. Basit, J. C. Chan, J. C. Mbanya *et al.*, "Idf diabetes atlas: Global, regional and country-level diabetes prevalence estimates for 2021 and projections for 2045," *Diabetes research and clinical practice*, 2022.

[3] C. Marling and R. Bunescu, "The ohiot1dm dataset for blood glucose level prediction: Update 2020," in *CEUR workshop proceedings*. NIH Public Access, 2020.

[4] D. C. Klonoff, C. Lias, R. Vigersky, W. Clarke, J. L. Parkes, D. B. Sacks, M. S. Kirkman, B. Kovatchev, and E. G. Panel, "The surveillance error grid," *Journal of diabetes science and technology*, 2014.

[5] 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)*, 2023.

[6] A. Chauhan, S. Bhattacharyya, and S. Vadivel, "Dqnas: Neural architecture search using reinforcement learning," *arXiv preprint arXiv:2301.06687*, 2023.

[7] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double q-learning," in *Proceedings of the AAAI conference on artificial intelligence*, 2016.

[8] R. M. Bergenstal, R. L. Gal, C. G. Connor, R. Gubitosi-Klug, D. Kruger, B. A. Olson, S. M. Willi, G. Aleppo, R. S. Weinstock, J. Wood *et al.*, "Racial differences in the relationship of glucose concentrations and hemoglobin a1c levels," *Annals of internal medicine*, 2017.