

NEURALCGMM: NEURAL CONTROL SYSTEM FOR CONTINUOUS GLUCOSE MONITORING AND MAINTENANCE

Azmine Toushik Wasi

Shahjalal University of Science and Technology, Bangladesh

azminetoushik.wasi@gmail.com

ABSTRACT

Precise glucose level monitoring is critical for people with diabetes to avoid serious complications. While there are several methods for continuous glucose level monitoring, research on maintenance devices is limited. To mitigate the gap, we provide a novel neural control system for continuous glucose monitoring and management that uses differential predictive control, *NeuralCGMM*. Our approach, led by a sophisticated neural policy and differentiable modeling, constantly adjusts insulin supply in real-time, thereby improving glucose level optimization in the body. This end-to-end method maximizes efficiency, providing personalized care and improved health outcomes, as confirmed by empirical evidence. Code and data are available at: <https://github.com/azminewasi/NeuralCGMM>.

1 INTRODUCTION

Continuous glucose monitoring and maintenance play a crucial role in diabetes management, offering real-time glucose level management that empowers informed decisions on diet, medication, and lifestyle, leading to improved glycemic control and reduced risk of hypoglycemia (Lee & Lupsa, 2021; Yu et al., 2023; Trief et al., 2016). Yet, there remains no suitable method capable of predicting, controlling, and automating the entire process efficiently. Differentiable Predictive Control (DPC) uses a differentiable programming-based policy gradient method to train a neural network to approximate an explicit Model Predictive Control (MPC) controller without the need for supervision from an expert controller (Cortez et al., 2022; 2023; Drgoňa et al., 2022). By employing differentiable system models, typically represented by ODEs, DPC adeptly adjusts to individual patient conditions, predicting and optimizing glucose levels.

In this work, we explore, define, and formulate the closed-loop GCM problem, and showcase a method (*NeuralCGMM*) to solve it effectively with DPC. Our framework for precise glucose-level management addresses a crucial gap in current healthcare systems. Formulated as a parametric optimal control problem, it minimizes tracking errors against desired references, incorporates tailored constraints for safe insulin injection, and provides personalized care.

2 PROBLEM FORMULATION

The ultimate goal of this work is to automate insulin delivery based on continuous glucose monitoring to maintain target glucose levels in individuals with diabetes. The control objective is to minimize the deviation of the glucose levels from the target values while optimizing insulin delivery to ensure stable and controlled blood glucose concentrations. The system is represented by a differentiable model that captures the dynamics of glucose-insulin interaction. The differentiable nature allows seamless integration with DPC. The control policy, denoted as $\pi(g_k, R)$, will determine insulin delivery actions at each time step based on the continuous glucose monitoring data g_k and the predictions $R = [r_k, \dots, r_{k+N}]$ over a predefined prediction horizon N . The parametric optimal control problem for glucose level maintenance is formulated as follows:

$$\underset{\theta}{\text{minimize}} \sum_{i=1}^m \left(\sum_{k=1}^{N-1} Q_g \cdot |g_k^i - r_k^i| + Q_N \cdot |g_N^i - r_N^i| + Q_u \cdot |u_k^i - u_k^{i-1}| \right) \quad (1)$$




Figure 1: (a) Continuous Glucose Monitoring and Maintenance System; (b) Model Data Flow; (c) Glucose Level Output. (Full-size figures are available in Appendix C.4)

subject to $g_{k+1}^i = \text{ODESolve}(f(g_k^i, u_k^i))$, $u_k^i = \pi_\theta(g_k^i, R^i)$, $g_0^i \sim \mathcal{P}_{g_0}$, $R^i \sim \mathcal{P}_R$, $u_k^i \in \mathcal{U}$, $g_k^i \in \mathcal{G}$.

The equations dictate the evolution of blood glucose levels, denoted as g_k^i , influenced by the differentiable system model $f(g_k^i, u_k^i)$, where u_k^i represents insulin delivery actions. The control policy π_θ determines insulin delivery based on glucose levels g_k^i and predicted references R^i . Initial glucose levels g_0^i follow a probability distribution \mathcal{P}_{g_0} , and references R^i follow \mathcal{P}_R . Constraints \mathcal{U} and \mathcal{G} ensure insulin delivery and glucose levels adhere to physiological limits. The formulation, incorporating weighting factors Q_g , Q_N , and Q_u , aims to optimize insulin delivery over a prediction horizon, addressing glucose control with consideration for safety and physiological factors. The optimization process involves tuning parameters θ using stochastic gradient descent.

3 NeuralCGMM METHOD AND EXPERIMENTS

Let us consider \mathcal{P}_{g_0} , \mathcal{P}_R and \mathcal{P}_D as distributions for initial conditions, required glucose levels, and system disturbances. For model-based policy optimization, we consider a discrete-time partially observable linear state space model (SSM) (Rangapuram et al., 2018) that characterizes the dynamics of a patient within a medical setting as a partially observable white-box system model (Dr̄goña et al., 2020). The model is represented by $g_{k+1} = Ag_k + Bu_k + Ed_k$; $y_k = Cg_k$. Here, g_k denotes the patient’s glucose level, u_k denotes control actions governing glucose regulation, and d_k encompasses system disturbances. y_k represents the measured variable—glucose flow regulation.

Next, we parameterize the control policy using deep neural networks, expressed as: $u_k = \pi_\theta(y_k, R, D)$. Here, y_k represents the insulin flow to be controlled, $R = \{y_{\min}, y_{\max}\}$ denotes the desired glucose level for the given insulin flow, and D corresponds to change in glucose level of the patient. With the partially observable system model and control policy in place, we formulate a differentiable closed-loop system model. In the closed-loop system, three penalty terms are integrated based on the objective function: control loss optimizes insulin regulation, a regularization loss deters aggressive changes in control actions to ensure patient safety, and a constraint loss limits insulin delivery rate and amount for feasibility, collectively optimizing the state variable g .

Integrating all components, we formulate a DPC problem to be optimized comprehensively over the distribution of training scenarios, as demonstrated in Figure 1 (b). We train the model for neural control policy using stochastic gradient descent. Details are discussed in Appendix B.

Experiments. To evaluate our model’s effectiveness, we conducted training using Neuromancer (Dr̄gona et al., 2023) with synthetic data, detailed in the Appendix. Experimental results, illustrated in Figure 1(c), showcase the model’s adept adaptation to evolving constraints during policy updates across 3000 simulation steps, highlighting the model’s responsiveness to changes in patient glucose levels and health dynamics. Details are discussed in Appendix C.

4 CONCLUSION

This work introduces a significant contribution by exploring, defining, and formulating the closed-loop GCM problem through the effective application of DPC. Our innovative framework, employing DPC for precise glucose management, addresses a critical gap in healthcare systems. Utilizing a sophisticated neural policy and prioritizing resource optimization, our system detects and automatically corrects glucose irregularities, enhancing patient support in critical healthcare situations.

ACKNOWLEDGEMENTS

We extend our sincere gratitude to all the reviewers for their invaluable feedback. Additionally, we would like to express our appreciation to Computational Intelligence and Operations Lab - CIOL for their support.

URM STATEMENT

Author Azmine Toushik Wasi meets the URM criteria of ICLR 2024 Tiny Papers Track.

REFERENCES

- Brandon Amos, Ivan Dario Jimenez Rodriguez, Jacob Sacks, Byron Boots, and J. Zico Kolter. Differentiable mpc for end-to-end planning and control, 2019. URL <https://arxiv.org/pdf/1810.13400.pdf>.
- Wenceslao Shaw Cortez, Jan Drgona, Aaron Tuor, Mahantesh Halappanavar, and Draguna Vrabie. Differentiable predictive control with safety guarantees: A control barrier function approach, 2022. URL <https://doi.org/10.48550/arXiv.2208.02319>.
- Wenceslao Shaw Cortez, Jan Drgona, Draguna Vrabie, and Mahantesh Halappanavar. Robust differentiable predictive control with safety guarantees: A predictive safety filter approach, 2023. URL <https://doi.org/10.48550/arXiv.2208.02319>.
- Jan Drgona, Aaron Tuor, James Koch, Madelyn Shapiro, and Draguna Vrabie. NeuroMANCER: Neural Modules with Adaptive Nonlinear Constraints and Efficient Regularizations. 2023.
- Ján Drgoňa, Damien Picard, and Lieve Helsen. Cloud-based implementation of white-box model predictive control for a geotabs office building: A field test demonstration. *Journal of Process Control*, 88:63–77, 2020. ISSN 0959-1524. doi: <https://doi.org/10.1016/j.jprocont.2020.02.007>. URL <https://www.sciencedirect.com/science/article/pii/S0959152419306857>.
- Ján Drgoňa, Karol Kiš, Aaron Tuor, Draguna Vrabie, and Martin Klaučo. Differentiable predictive control: Deep learning alternative to explicit model predictive control for unknown nonlinear systems. *Journal of Process Control*, 116:80–92, 2022. ISSN 0959-1524. doi: <https://doi.org/10.1016/j.jprocont.2022.06.001>. URL <https://www.sciencedirect.com/science/article/pii/S0959152422000981>.
- Grace S Lee and Beatrice C Lupsa. Continuous glucose monitoring for the internist. *Med. Clin. North Am.*, 105(6):967–982, Nov 2021. doi: [10.1016/j.mcna.2021.06.004](https://doi.org/10.1016/j.mcna.2021.06.004). URL <https://doi.org/10.1016/j.mcna.2021.06.004>.
- Rui Ma, Ruomei Shao, Xuyao An, Qichun Zhang, and Shuqing Sun. Recent advancements in noninvasive glucose monitoring and closed-loop management systems for diabetes. *J. Mater. Chem. B*, 10:5537–5555, 2022. doi: [10.1039/D2TB00749E](https://doi.org/10.1039/D2TB00749E). URL <http://dx.doi.org/10.1039/D2TB00749E>.
- Eden M. Miller. Using Continuous Glucose Monitoring in Clinical Practice. *Clinical Diabetes*, 38(5):429–438, 12 2020. ISSN 0891-8929. doi: [10.2337/cd20-0043](https://doi.org/10.2337/cd20-0043). URL <https://doi.org/10.2337/cd20-0043>.
- Alex Oshin and Evangelos A. Theodorou. Differentiable robust model predictive control, 2023. URL <https://arxiv.org/pdf/2308.08426.pdf>.
- Syama Sundar Rangapuram, Matthias W Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski. Deep state space models for time series forecasting. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/5cf68969fb67aa6082363a6d4e6468e2-Paper.pdf.

Oliver Schubert-Olesen, Jens Kröger, Thorsten Siegmund, Ulrike Thurm, and Martin Halle. Continuous glucose monitoring and physical activity. *International Journal of Environmental Research and Public Health*, 19(19), 2022. ISSN 1660-4601. doi: 10.3390/ijerph191912296. URL <https://www.mdpi.com/1660-4601/19/19/12296>.

Paula M. Trief, Donald Cibula, Elaine Rodriguez, Bridget Akel, and Ruth S. Weinstock. Incorrect Insulin Administration: A Problem That Warrants Attention. *Clinical Diabetes*, 34(1):25–33, 01 2016. ISSN 0891-8929. doi: 10.2337/diaclin.34.1.25. URL <https://doi.org/10.2337/diaclin.34.1.25>.

Jin Yu, Jae-Hyoung Cho, and Seung-Hwan Lee. The era of continuous glucose monitoring and its expanded role in type 2 diabetes. *Journal of Diabetes Investigation*, 14(7):841–843, 2023. doi: <https://doi.org/10.1111/jdi.14028>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/jdi.14028>.

A RELATED WORKS AND MOTIVATION

A.1 CONTINUOUS GLUCOSE MONITORING AND MAINTENANCE

Continuous Glucose Monitoring (CGM) technology, introduced in 1999, has emerged as a transformative tool in diabetes management. Offering real-time and predictive glycemic data, CGM systems contribute significantly to detecting trends, identifying asymptomatic events, and assessing glycemic variability. Enhanced frequency of glucose monitoring correlates with reduced hypoglycemia and increased time in range, leading to improved A1C levels. The comprehensive insights provided by CGM data analysis enable targeted treatment interventions, such as preventing hypoglycemia, optimizing glycemic control at specific times, and enhancing overall performance (Miller, 2020).

Recent advancements in CGM systems go beyond traditional blood glucose monitoring, offering robust data for effective diabetes management. Schubert-Olesen et al. outlines practical recommendations for incorporating CGM into everyday physical activities, emphasizing the pivotal role of CGM in enhancing patient engagement. Additionally, Ma et al. underscores the significance of closed-loop management systems, integrating electrochemical sensing of glucose and noninvasive monitoring technologies. These advancements not only underscore the potential of CGM in improving glucose control, especially in patients with higher initial A1C levels, but also emphasize the importance of consistent device usage for optimal benefits, setting a foundation for the future of diabetes care.

Despite notable advancements in glucose monitoring technologies, there exists a notable gap in addressing maintenance aspects, particularly the development of a system capable of autonomously administering the required insulin dosage to patients as needed. The complexity of such a system lies in its integration of both an accurate prediction policy and an optimization policy, with constraints aimed at minimizing costs and risks while simultaneously ensuring improved patient health outcomes.

A.2 DIFFERENTIABLE PREDICTIVE CONTROL

Differentiable Predictive Control (DPC) stands as an innovative methodology for addressing the challenges of model predictive control (MPC) in complex systems. By learning explicit neural control laws offline, DPC effectively mitigates the computational demands of online MPC. It achieves this by incorporating state and input constraints into the loss function through penalty functions and aggregating them with the MPC cost function. The resulting neural network control policy is trained offline using stochastic gradient descent, leveraging automatic differentiation of MPC problem cost functions and constraints. Demonstrating high performance with low computational resources, DPC has been successfully implemented in various applications, offering a promising avenue for enhancing control system efficiency and applicability (Drgoña et al., 2022; Cortez et al., 2022; Oshin & Theodorou, 2023).

Recent advancements in DPC extend its utility by introducing it as a differentiable policy class for reinforcement learning in continuous state and action spaces. By leveraging KKT conditions and convex approximation, researchers have enabled the differentiation through MPC, allowing for

end-to-end learning of the cost and dynamics of a controller. Notably, DPC exhibits superior data efficiency in comparison to generic neural networks, particularly evident in experiments involving pendulum and cartpole domains (Amos et al., 2019). The approach marks a departure from traditional system identification methods, showcasing DPC’s potential to outperform existing techniques and streamline learning in scenarios where expert guidance may be unavailable.

Recognizing DPC’s capability in managing intricate systems, we plan to address the maintenance challenges of current CGM systems by developing an advanced closed-loop system presented in this paper. This solution aims to enhance the effectiveness, efficiency, and optimization of continuous glucose monitoring and management using a neural network control policy.

B METHODOLOGY IN DEPTH

B.1 DIFFERENTIABLE PREDICTIVE CONTROL

DPC, a deep learning-driven substitute for MPC in handling unknown nonlinear systems, utilizes deep learning to grasp system dynamics from data, providing adaptability to such systems. DPC excels in managing complex real-world scenarios due to its end-to-end optimization, flexibility, and constraint support. Its ability to seamlessly handle time-varying references and constraints, along with reduced computational complexity, establishes DPC as a more efficient and versatile alternative to conventional MPC in control applications.

DPC is an algorithm rooted in model-based policy optimization. It leverages the differentiability inherent in a broad range of model representations for dynamic systems, encompassing differential equations, state-space models, and diverse neural network architectures. In the DPC framework, a differentiable closed-loop system is created, consisting of a neural control policy and a system dynamics model. The optimization process involves refining this system by utilizing parametric control objectives as intrinsic reward signals. These rewards are evaluated across a sampled distribution of the problem parameters, enabling the algorithm to iteratively improve the performance of the control policy and model (Drgoňa et al., 2022). A sample diagram is added here from the original DPC work.

The DPC problem can be formally expressed as a parametric optimal control problem, where a parametric control policy, denoted as $\pi_\theta(g(t), \xi(t))$ and parameterized by trainable weights W , is employed for a continuous-time dynamical system. The system is described by the differential equation:

$$\frac{dg(t)}{dt} = f(g(t), u(t)) \quad (2)$$

Here, $g(t)$ represents the time-varying state, $u(t)$ denotes the system control inputs, and f characterizes the state transition dynamics. Alternatively, in discrete-time form, obtained through methods such as ODE solvers or state-space models, the system evolves as:

$$g_{k+1} = f(g_k, u_k) \quad (3)$$

The objective is to optimize the parametric control policy to determine the control inputs, u , in order to minimize a certain cost or objective function associated with the system’s behavior. The optimization is performed over the parameter space θ to enhance the policy’s performance in guiding the system dynamics. Using, this, we formulate the DPC problem as a following parametric optimal control problem presented in Equation 1.

$$\underset{\theta}{\text{minimize}} \sum_{i=1}^m \left(\sum_{k=1}^{N-1} Q_g \cdot |g_k^i - r_k^i| + Q_N \cdot |g_N^i - r_N^i| + Q_u \cdot |u_k^i - u_k^{i-1}| \right) \quad (4)$$

subject to:

$$\begin{aligned} g_{k+1}^i &= \text{ODESolve}(f(g_k^i, u_k^i)), \\ u_k^i &= \pi_\theta(g_k^i, R^i), \\ g_0^i &\sim \mathcal{P}_{g_0}, \\ R^i &\sim \mathcal{P}_R, \\ u_k^i &\in \mathcal{U}, \\ g_k^i &\in \mathcal{G}. \end{aligned}$$

The primary benefit of formulating the DPC problem with differentiable closed-loop dynamics models, control objective functions, and constraints lies in the ability to employ automatic differentiation, specifically backpropagation through time. This approach facilitates the direct computation of policy gradients. By expressing the problem as a computational graph and leveraging the chain rule, we can efficiently compute the gradients of the loss function with respect to the policy parameters W .

To obtain the gradients of the objective function with respect to the policy parameters W , we can apply the chain rule for partial differentiation. Let’s denote the objective function as J :

$$J(\theta) = \sum_{i=1}^m \left(\sum_{k=1}^{N-1} Q_g \cdot |g_k^i - r_k^i| + Q_N \cdot |g_N^i - r_N^i| + Q_u \cdot |u_k^i - u_k^{i-1}| \right) \quad (5)$$

The chain rule for partial differentiation of J with respect to the policy parameters W is expressed as:

$$\frac{\partial J}{\partial W} = \sum_{i=1}^m \sum_{k=1}^{N-1} \left(\frac{\partial J}{\partial g_k^i} \cdot \frac{\partial g_k^i}{\partial W} + \frac{\partial J}{\partial u_k^i} \cdot \frac{\partial u_k^i}{\partial W} \right) \quad (6)$$

Here, $\frac{\partial J}{\partial g_k^i}$ and $\frac{\partial J}{\partial u_k^i}$ are the partial derivatives of the objective function with respect to the state and control inputs at time step k , respectively. The terms $\frac{\partial g_k^i}{\partial W}$ and $\frac{\partial u_k^i}{\partial W}$ represent the gradients of the state and control inputs with respect to the policy parameters W .

B.2 WHITE-BOX SYSTEM MODEL WITH CONTROL POLICY

As mentioned in Section 3, the control policy is parameterized using deep neural networks (MLP), specifically expressed as $u_k = \pi_\theta(y_k, R, D)$. In this formulation, y_k represents the glucose level to be regulated, $R = \{y_{\min}, y_{\max}\}$ denotes the desired glucose levels of the patient, and D accounts for observed disturbances, encompassing changes in the patient due to physiological factors such as eating or drinking. The neural network structure is instantiated as an MLP with bounds, characterized by an input size comprising the current glucose level, desired glucose levels, and observed disturbances. It outputs the control input u_k . The overall neural network policy, is constructed as a node in the computational graph, mapping inputs to the corresponding control output as follows in Figure 2.

C EXPERIMENTS AND IMPLEMENTATION DETAILS

C.1 PARAMETERS

All experimental information is provided in this table 1.

C.2 DATASET GENERATION

The data generation process involves creating training and development datasets for a dynamic system. For each scenario, initial conditions are sampled from the system, and a prediction horizon of 100 steps is defined. The optimal is sampled from a uniform distribution between 18 and 20. Disturbance trajectories are generated using the system’s simulation model. The data is organized into batches for efficient training, with 64 sampled scenarios for each batch. Finally, dataloader




Figure 2: Close Loop System Model.

Table 1: Parameter Settings for the System

NAME	VALUE
Desired Glucose Level, g_{\min_range}	(12., 18.)
Prediction Horizon, n_{steps}	100
Samples, $n_{samples}$	2000
Batch Size	64
Control Model	MLP
MLP Specification	2 hidden layers with 32 units in each
MLP Activation	GELU
Control Loss Hyperparameter	$Q_g = 0.01$
Regularization Loss Hyperparameter	$Q_N = 0.1$
Constraint Loss Hyperparameter	$Q_u = 0.02$
Trainer Parameters: Epochs	200
Trainer Parameters: Warmup	50
Trainer Parameters: Optimizer	AdamW
Trainer Parameters: Learning Rate	0.001
Trainer Parameters: Early Stopping	No update in 5 steps

instances are created for both training and development datasets, facilitating the training of the DPC algorithm on the dynamic system.

C.3 TEST

For testing, we randomly generated a test set of 3000 datapoints as described in Section C.2, and generated predictions using the model.

In Figure 3, we can see the output and model parameters for different values in the synthetic test set. We can see in the y chart that the control policy is trying to keep the value within limits, though there are several changes in the system. u is also within the constraints. Also, the d chart shows that the model is capturing the disturbances nicely.

The best model and test data samples, along with the codes, are included in the supplementary material for result reproduction purposes.




Figure 3: Test Output Visualization.

C.4 FIGURES

Here we include a zoomed version of the two charts in Figure 1. The another is already added in Figure 3 chart 1.

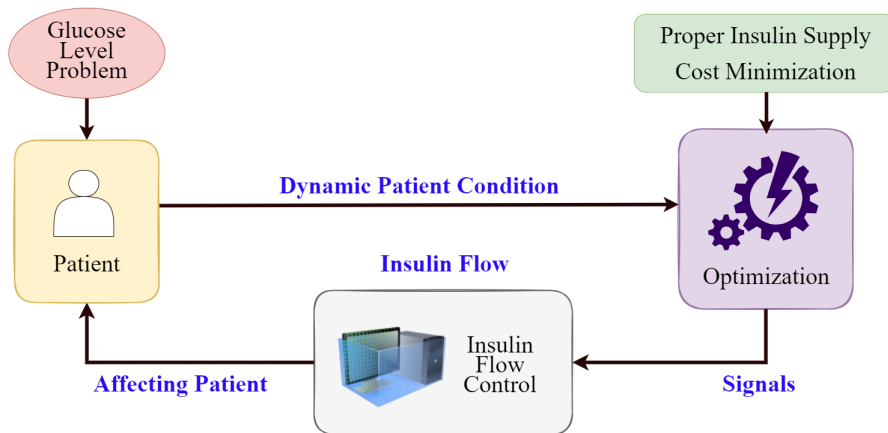


Figure 4: Continuous Glucose Monitoring and Maintenance System




Figure 5: Model Data Flow for Continuous Glucose Monitoring and Maintenance System

D DISCUSSION

This study establishes a theoretical foundation for monitoring and maintaining glucose levels through differential predictive control. We feel that it is a scalable procedure that requires in-depth research and development with clinical objectives in order to develop a practical healthcare solution.

Study Limitations due to Synthetic Data. We acknowledge that synthetically generated data has limitations and might not accurately represent the intricacies of real-world situations. It is important to point out that this relates to a clinical procedure in which errors could result in patient deaths. As the research on active glucose level maintenance is still limited, we are unable to test it with real-world data. Using synthetically generated data, this work provides a foundational framework for further exploration and clinical studies.

Scalability. Usual healthcare devices face scalability challenges in diverse clinical settings due to varying patient conditions and limitations. However, DPC offers a promising solution by dynamically adapting to patient dynamics, mitigating the impact of clinical variations. Its ability to handle personalized constraints and optimize insulin delivery makes it a scalable and adaptable approach for addressing diverse clinical limitations in our parametric optimal control framework. Effective implementation of this approach requires extensive research, rigorous testing, meticulous calibration, and thorough clinical trials to setup and analyze different parameters and factors, ultimately leading to the development of a practical product that addresses glucose level management effectively.

Limitations. There are certain shortcomings and problems that may arise that should be considered. The extensive calibration required for the practical application of our approach can be a practical obstacle in real-world healthcare settings, as can the complexity of conducting rigorous clinical testing during the development stage. These tasks can be costly and time-consuming, which may hinder the broad adoption of our approach. Lastly, unexpected external factors and technological limitations may affect the precision and dependability of glucose level predictions, which we believe is a problem for almost all devices.

Through this work, we hope to pave the way for future research initiatives, ultimately contributing to the development of a refined device that holds the potential to positively impact the lives of millions worldwide.