

Food Recommendation Using Machine Learning for Physical Activities in Patients with Type 1 Diabetes

Phuong Ngo¹, Maryam Tayefi², Anne Torill Nordsletta³ and Fred Godtliebsen⁴

¹Norwegian Centre for E-health Research, Norway, phuong.dinh.ngo@ehealthresearch.no

²Norwegian Centre for E-health Research, Norway, maryam.tayefi@ehealthresearch.no

³Norwegian Centre for E-health Research, Norway, anne.torill.nordsletta@ehealthresearch.no

⁴UiT The Arctic University of Norway, Tromsø, Norway, fred.godtliebsen@uit.no

Abstract

Physical activities have a significant impact on blood glucose homeostasis of patients with type 1 diabetes. Regular physical exercise provides many proven health benefits and is recommended as part of a healthy lifestyle. However, one of the main side effects of physical activities is hypoglycemia (low blood glucose). Fear of hypoglycemia generally leads to the patients not participating in physical activities. This paper shows a proof of concept that machine learning can be used to create a personalized food recommendation system for patients with type 1 diabetes. Machine learning algorithms were designed to improve glycemic control and reduce the overcompensation of carbohydrate. First, a personalized model based on feedforward neural networks is developed to predict the blood glucose outcome during and after physical activities. Based on the personalized model and reinforcement learning, optimal food intakes will be recommended to the patient. Simulation results show that the proposed methodology has successfully maintained the blood glucose in the healthy range on a type 1 diabetes simulator during physical activities.

Keywords

Type 1 diabetes, physical activities, feedforward neural network, reinforcement learning.

1 INTRODUCTION

Type 1 diabetes is a chronic disease characterized by the lack of insulin secretion due to the autoimmune destruction of pancreatic beta cells. This results in an uncontrolled increase of blood glucose level. High blood glucose (BG) level can lead to complications and eventually failure of various organs in the body. On the other hand, low BG level (hypoglycemia) is an acute complication of diabetes. Hypoglycemia is defined when the BG is dropped to less than 3.9 mmol/l [1]. Hypoglycemia causes symptoms from increased heart rate to mental confusion, and unconsciousness. Repeated episodes of hypoglycemia can also lead to brain damage. For many patients with diabetes, the hypoglycemia symptoms can be hard to detect due to a phenomenon, called hypoglycemic unawareness. Hypoglycemia unawareness is very dangerous as BG level may approach extremely low before any symptoms are perceived [2]–[4].

Regular physical exercises have many health benefits and are therefore widely recommended for patients with type 1 diabetes. However, exercises alter significantly glucose homeostasis in patients with type 1 diabetes [5], [6]. Physical activities increase glucose uptake by muscles leading to a drop in BG concentration, which can reach the hypoglycemic values. Besides, increased insulin sensitivity effects are long-lasting after physical activities and have many negative impacts on daily activities of patients.

Since an automated solution in controlling blood glucose can bring tremendous benefit for patients with type 1

diabetes, various studies have been conducted to design algorithms for this purpose. For example, Marchetti *et al.* [7] derived a proportional integral derivative controller for BG control. Soylu *et al.* [8] proposed a Mamdani type fuzzy control strategy for exogenous insulin infusion. However, the glucose kinetics process is complex [9] and depends on many factors such as food intakes, active insulin, physical activities, stress, and hormone changes. Furthermore, many of the techniques for BG control are difficult to be implemented since they either require extensive manual tuning for adapting to individual patients or assume that an accurate mathematical model of the patient BG dynamics is available.

Recently, machine learning algorithms have been widely used since they are able to learn and gain intelligence by utilizing a large amount of available data generated by the development of new technologies. For example, artificial neural networks (ANN) is an effective method that imitates how a nervous system works in a simple way and can be used for obtaining a personalized model of BG activities. Reinforcement learning (RL) is also a suitable machine learning tool for BG control. RL was developed and studied in control theory [10], game theory [11], information theory [12] and applied in many other applications including diabetes [13]–[15]. Through a series of experiments, Fox and Wiens [16] compare the performance of different RL approaches to non-RL approaches and concluded that RL is a promising tool for improving blood glucose for individuals with type 1 diabetes. In this paper, novel, safe-for-patients machine learning techniques will be studied and developed in order

to provide an estimation of food for patients with type 1 diabetes.

2 MACHINE LEARNING ALGORITHMS FOR FOOD RECOMMENDATION TO PATIENTS WITH TYPE-1 DIABETES

Depending on the length of physical activities, two alternatives for food recommendation can be provided to patients. For short physical activities, the system will recommend patients to eat only at the beginning of the exercise. The amount of carbohydrate (CHO) is recommended based on the prediction of the BG outcome from the feedforward neural network described in this section. For long physical activities, it is necessary to distribute food intake during the activities to keep the BG stable. RL is used to estimate the optimal distribution of food intake during the exercise for this purpose.

2.1 Model-Based Recommendation for Short Physical Activities Using Feedforward Neural Networks

A feedforward neural network (FFNN) is a type of ANN which is constructed by neurons organized into layers. The network can be used to estimate the blood glucose outcome from the information of the food that the patient consumes, the amount of physical activity and other factors. A structure of a simple FFNN demonstrated in this paper can be found in Figure 1.




Figure 1 Diagram of the FFNN for estimating blood glucose outcome during physical activities

In this simple network, the inputs are the amount of CHO in food consumed by the patient before doing exercises and the average heart rate expected during the physical activity. The output is the BG outcome, which is represented by a score assigned for the average BG level during and after physical activities (Table 1). The score (ranged from -10 to 10) is designed such that it is high when the BG is closer to the healthy level and low when the BG is further away from the healthy value.

In FFNN, information flows from inputs through the hidden layer towards the output. Each node in the hidden layer is a rectifying linear unit function (ReLU) that mimics how the electrical impulse is fired from one neuron to another in the human brain. The output signal from each node can be represented mathematically as follows:

$$a_j^l = \sigma \left(\sum_k w_{jk}^l a_k^{l-1} + b_j^l \right) \quad (1)$$

where σ is the activation function, a_j^l is the output value of node j in layer l . The notation w_{jk}^l is the weight of the connection from node k in layer $l-1$ to node j in layer l . b_j^l

denotes the bias of node j in layer l . The task of training a neural network is to find the optimal set of w_{jk}^l and b_j^l such that the following cost function is minimized:

$$MSE = \frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2 \quad (2)$$

where N is the number of training samples, \hat{y}_i is the predicted value and y_i is the actual blood glucose score for each training sample i .

Upon trained, the optimal amount of CHO in food A^* that the patient should consume before the exercise can be obtained from the FFNN as follows:

$$A^* = \operatorname{argmax}_a f(a) \quad (3)$$

where $f(a)$ is the mapping between the amount of carbohydrate in food and the blood glucose score.

2.2 Food Recommendation for Long Physical Activities Using Reinforcement Learning

The principle of RL is based on the interaction between a decision-making agent and its environment [17]. At certain times during the physical activity, the recommendation system evaluates the current BG condition and makes a recommendation of the amount of food that the patient should eat. The evolution of the BG as the results of the action by the patient determines whether the patient gets a positive or negative reinforcement (blood glucose score). Mathematically, the RL framework for recommending food for patients with type 1 diabetes during physical activities consists of the following elements:

- The state $s = S_t$ defines the condition of the patient at time t . It includes historical values of the blood glucose levels and the information about the physical activity intensity.
- The action $a = A_t$ (eg: type and quantity of food) that follows a policy $\pi(s, a)$. A policy is a mapping between the current condition of the patient and the probabilities of selecting each possible action.
- The score/reward $r = R_{t+1}$, which is the result (consequence) of action A_t at the state S_t .

The objective of the algorithm is to keep the BG level within the healthy level as much as possible during the physical activity. Hence it will search for an optimal policy that will maximize the accumulation of score/reward throughout the exercise. The accumulation of score/reward at state s when taking action a is defined as the action value function:

$$q_\pi(s, a) = \mathbb{E}_\pi \left\{ \sum_{k=0}^{\infty} \gamma^k R_{\{t+k+1\}} \mid S_t = s, A_t = a \right\} \quad (4)$$

With \mathcal{S} as the set of all possible states and \mathcal{A} as the set of all possible actions, the ϵ -greedy policy obtained from the action value function is defined as follows:

$$\pi(a, s) = \begin{cases} 1 - \epsilon + \epsilon/|\mathcal{A}| & \text{if } a = A^* \\ \epsilon/|\mathcal{A}| & \text{if } a \neq A^* \end{cases} \quad (5)$$

for all $s \in \mathcal{S}$, $a \in \mathcal{A}$, and A^* is the optimal food action:

$$A^* = \operatorname{argmax}_a Q(S_t, a).$$

The algorithm for controlling the BG during physical activity using RL can be summarized in Table 1.

3 RESULTS AND DISCUSSION

In order to demonstrate how the algorithms work, we have built a glucose kinetics simulator based on the physical activity model suggested by Breton [18] and part of the Hovorka's model [19] which describes the CHO absorption process in the body. The mathematical description of the simulator can be found in the Appendix.

Initialize the estimated action value function $Q(s, a)$ for all $s \in \mathcal{S}$ and $a \in \mathcal{A}$.

Obtain the ϵ -greedy policy from the initial estimated action value function.

For each exercise do:

For each break t during the exercise do:

- Suggest an amount of food A_t based on the current policy $\pi(s, a)$.
- Collect the dataset: $S_{t-1}, A_{t-1}, S_t, A_t$.
- Update the current policy from the estimated action value function.

end

end

Table 1 Reinforcement-learning algorithm for food recommendation during physical activity for patients with type 1 diabetes.

BG level	Score/reward
BG < 3.9 mmol/L	-10
3.9 mmol/L ≤ BG < 4.2 mmol/L	-3
4.2 mg/dl ≤ BG < 5.6 mmol/L	10
5.6 mmol/L ≤ BG < 7.2 mmol/L	5
7.2 mmol/L ≤ BG < 10.0 mmol/L	-1
10.0 mmol/L ≤ BG < 15.6 mmol/L	-5
BG ≥ 15.6 mmol/L	-8

Table 2 Score/reward for different BG levels.

3.1 Short Physical Activities

For short physical activities, the recommendation is based on the FNN and is given at the beginning of each exercise. Training data for the FFNN was obtained by repeated simulations from the BG simulator under scenarios that a patient with type 1 diabetes performs physical exercises with different intensities and consumed different amount of food. The duration of physical activities is set to be constant at 30 minutes and the patient always eat at 15 minutes before the exercise starts. The outcome of each exercise is evaluated by measuring the average scores of the BG (defined in Table 2) over the course of three hours starting at 15 minutes before the exercise. The BG is sampled every 5 minutes during the simulations, which is similar to the sampling time of many continuous glucose monitoring devices. Physical intensities are represented by heart rate values in the

simulator. It is also assumed that during short physical activities, the heart rates are constant.

The result of the trained neural networks is shown in Figure 2 as a mapping from food amount and heart rate to the blood glucose outcome during each exercise. Based on this mapping, the optimal amount of food was calculated and given to the simulator.

A test scenario was carried out in which the average heart of the exercise is 100 bpm and the patient is provided with four different choices of food intakes before the exercise: 0, 10, 20, 30 and 40 grams of CHO. Figure 3 shows the comparison of the blood glucose responses within three hours when the patient does not eat anything, follows the recommendation from the algorithm (10 grams of CHO) or have the highest CHO portion (40 grams) at 15 minutes before the exercise. The result shows that by following the recommended food obtained based on the FFNN, the BG of the patient stays well within the healthy level for the duration of the simulation.




Figure 2 Estimation of the blood glucose score using feedforward neural network (blue dots represent training data).

3.2 Long Physical Activities

In the long physical activity scenario, a patient performs an interval exercise with the heart-rate profile as shown in Figure 4. Three cases were simulated in this scenario. In case 1, the patient consumes food with the recommended CHO provided by the FFNN. In case 2, the patient is recommended by the RL algorithm the amount of food to eat at the beginning and at every 20 minutes during the exercise. The choices of actions suggested by the recommendation system include: eat nothing, one portion or two portions of food. Each portion of food contains 10 grams of CHO. In other words, at every 20 minutes during the exercise, the algorithm suggests the patient how much he or she should eat to keep the blood glucose level in the healthy level based on the blood glucose data (sampled at every 5 minutes) in the previous 20 minutes. After each time the patient eats, a reward or score is given based on the blood glucose responses in the next 20 minutes. The value of the reward is assigned based on different blood glucose levels and is provided in Table 2. In case 3, the patient does not eat any food at the beginning and throughout the physical activity.

Figure 5 shows the BG responses proposed by the RL

algorithm and the FFNN based recommendation in the same simulation scenario. Figure 6 shows the amount of food (CHO) recommended by the RL and the FFNN. From Figure 5, it can be seen that the BG has been well regulated. No hypoglycemia occurs in both scenarios. However, by spreading out and calculating the optimal food consumption throughout the period of long exercises as shown in Figure 6, RL has better performances compared to the FFNN.




Figure 3 Blood glucose responses with different amount of CHO during a short physical activity (lighter shades indicate more desirable BG levels).




Figure 4 Heart rate during the exercise in the long physical activity scenario




Figure 5 BG responses during the interval exercise in our simulations (Case 1: Food consumed at the beginning of physical activity using FFNN, Case 2: Food consumed throughout physical activity using RL, Case 3: No food consumption).

4 CONCLUSION

This paper provides two algorithms that can be used in a food recommendation system for patients with type 1 diabetes: the model-based method based on feedforward neural networks and the reinforcement learning method. Simulation results show that the feedforward neural network based method is suitable for the scenario when the length of the exercise is short and data from past physical activities are available. However, reinforcement learning performs better in situations where physical activities are long and food intakes can be spread out during exercises.




Figure 6 Food amounts recommended by the feedforward neural network (case 1) and the reinforcement learning (case 2).

5 REFERENCES

- [1] Seaquist, E. R. *et al.*, "Hypoglycemia and Diabetes: A Report of a Workgroup of the American Diabetes Association and The Endocrine Society," *Diabetes Care*, vol. 36, no. 5, pp. 1384–1395, 2013.
- [2] de Galan, B. E., Schouwenberg, B. J. J. W., Tack, C. J., and Smits, P., "Pathophysiology and management of recurrent hypoglycaemia and hypoglycaemia unawareness in diabetes.," *The Netherlands journal of medicine*, vol. 64, no. 8, pp. 269–79, 2006.
- [3] Czyzewska, K., Czerniawska, E., and Szadkowska, A., "Prevalence of hypoglycemia unawareness in patients with type 1 diabetes," in *Abstracts for the 38th Annual Meeting of the International Society for Pediatric and Adolescent Diabetes (ISPAD)*, 2000, vol. 13.
- [4] Schopman, J. E., Geddes, J., and Frier, B. M., "Prevalence of impaired awareness of hypoglycaemia and frequency of hypoglycaemia in insulin-treated type 2 diabetes.," *Diabetes research and clinical practice*, vol. 87, no. 1, pp. 64–8, 2010.
- [5] Riddell, M. and Perkins, B. A., "Exercise and glucose metabolism in persons with diabetes mellitus: Perspectives on the role for continuous glucose monitoring," *Journal of Diabetes Science and Technology*, vol. 3, no. 4, pp. 914–923, 2009.
- [6] Camacho, R. C., Galassetti, P., Davis, S. N., and Wasserman, D. H., "Glucoregulation during and after exercise in health and insulin-dependent diabetes.," *Exercise and sport sciences reviews*, vol. 33, no. 1, pp. 17–23, 2005.

- [7] Marchetti, G., Barolo, M., Jovanovic, L., Zisser, H., and Seborg, D. E., “An improved PID switching control strategy for type 1 diabetes,” *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 3, pp. 857–865, 2008.
- [8] Soyly, S., Danisman, K., Sacu, I. E., and Alci, M., “Closed-loop control of blood glucose level in type-1 diabetics: A simulation study,” in *International Conference on Electrical and Electronics Engineering (ELECO)*, 2013, pp. 371–375.
- [9] Wang, Q. *et al.*, “Personalized state-space modeling of glucose dynamics for type 1 diabetes using continuously monitored glucose, insulin dose, and meal intake: an extended Kalman filter approach,” *Journal of Diabetes Science and Technology*, vol. 8, no. 2, pp. 331–345, 2014.
- [10] Vrabie, D., Kyriakos G. Vamvoudakis, and Frank L. Lewis, *Optimal Adaptive Control and Differential Games by Reinforcement Learning Principles*, 1st ed., vol. 81. London: Institution of Engineering and Technology, 2012.
- [11] Lanctot, M. *et al.*, “A unified game-theoretic approach to multiagent reinforcement learning,” in *31st International Conference on Neural Information Processing Systems*, 2017, pp. 4191–4204.
- [12] Leibfried, F., Grau-Moya, J., and Bou-Ammar, H., “An Information-Theoretic Optimality Principle for Deep Reinforcement Learning,” in *Deep Reinforcement Learning Workshop NIPS*, 2018.
- [13] Bothe, M. K. *et al.*, “The use of reinforcement learning algorithms to meet the challenges of an artificial pancreas,” *Expert Review of Medical Devices*, vol. 10, no. 5, pp. 661–673, 2013.
- [14] De Paula, M., Ávila, L. O., and Martínez, E. C., “Controlling blood glucose variability under uncertainty using reinforcement learning and Gaussian processes,” *Applied Soft Computing Journal*, vol. 35, pp. 310–332, 2015.
- [15] Ngo, P. D., Wei, S., Holubova, A., Muzik, J., and Godtliebsen, F., “Reinforcement-learning optimal control for type-1 diabetes,” in *2018 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*, 2018, pp. 333–336.
- [16] Fox, I. and Wiens, J., “Reinforcement Learning for Blood Glucose Control: Challenges and Opportunities,” in *Workshop in the 36th International Conference on Machine Learning*, 2019.
- [17] Sutton, R. and Barto, A., *Reinforcement Learning: An Introduction*, 1st ed. Cambridge, MA: MIT Press, 2018.
- [18] Breton, M. D., “Physical Activity—The Major Unaccounted Impediment to Closed Loop Control,” *Journal of Diabetes Science and Technology*, vol. 2, no. 1, pp. 169–174, 2008.
- [19] Hovorka, R. *et al.*, “Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes,” *Physiological Measurement*, vol. 25, no. 4, pp. 905–20, 2004.

6 APPENDIX

The blood glucose simulator used in the paper was constructed based on the physical activity model suggested by Breton [18] and part of the Hovorka model [19]:

$$\begin{aligned}\frac{dD_1(t)}{dt} &= A_G D(t) - \frac{D_1(t)}{\tau_D} \\ \frac{dD_2(t)}{dt} &= \frac{D_1(t)}{\tau_D} - \frac{D_2(t)}{\tau_D} \\ \frac{dg}{dt} &= -p_1 g(t) + \frac{D_2(t)}{\tau_D} - \chi(t)g(t) - \beta Y Q_1 \\ \frac{d\chi}{dt} &= -p_2 \chi(t) + p_3 V(i(t) - i_b(t))\end{aligned}$$

Descriptions of the variables and parameter values can be found in Table 3 and Table 4.

Parameters	Value
p_1 Glucose effectiveness	0.2 min^{-1}
p_2 Insulin sensitivity	0.028 min^{-1}
p_3 Insulin rate of clearance	10^{-4} min^{-1}
A_G CHO bioavailability	0.8 min^{-1}
τ_D Glucose absorption constant	10 min
V Plasma volume	2730 g
$i_b(t)$ Initial basal rate	$7.326 \mu\text{IU}/(\text{ml}\cdot\text{min})$

Table 3 Parameters of the blood glucose simulator.

Variable	Unit
D Amount of CHO intake	mmol/min
D_1 Glucose in compartment 1	Mmol
D_2 Glucose in compartment 2	Mmol
$g(t)$ Plasma glucose concentration	mmol/l
$\chi(t)$ Interstitial insulin activity	min ⁻¹
$i(t)$ Plasma insulin concentration	$\mu\text{IU}/\text{ml}$

Table 4 Variables of the blood glucose kinetics model.

