Evaluation of Blood Glucose Level Control in Type 1 Diabetic Patients Using Deep Reinforcement Learning

Phuwadol Viroonluecha (✉ phuwadol.viroonluecha@upct.es)

Universidad Politecnica de Cartagena (UPCT)

Esteban Egea-Lopez

Universidad Politecnica de Cartagena (UPCT)

Jose Santa

Universidad Politecnica de Cartagena (UPCT)

Research Article

Keywords: Diabetes mellitus, insulin pump, glucose levels, simple reward function, safe glycemic state and low risk.

Posted Date: December 1st, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-1095721/v1>

License: This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at PLOS ONE on September 13th, 2022. See the published version at <https://doi.org/10.1371/journal.pone.0274608>.

Evaluation of Blood Glucose Level Control in Type 1 Diabetic Patients Using Deep Reinforcement Learning

Phuwadol Viroonluecha^{1,*,+}, Esteban Egea-Lopez^{1,+}, and Jose Santa^{2,+}

¹Universidad Politecnica de Cartagena (UPCT), Department of Information Technologies and Communications, Cartagena, 30202, Spain

²Universidad Politecnica de Cartagena (UPCT), Department of Electronics, Computer Technology and Projects, Cartagena, 30202, Spain

*phuwadol.viroonluecha@upct.es

+these authors contributed equally to this work

ABSTRACT

Diabetes mellitus is a disease associated with abnormally high levels of blood glucose due to lack of insulin. Combining an insulin pump and continuous glucose monitor with a control algorithm to deliver insulin is an alternative to patient self-management of insulin doses to control blood glucose levels in diabetes mellitus patients. In this work we propose a closed-loop control for blood glucose levels based on deep reinforcement learning. We describe the initial evaluation of several alternatives conducted on a realistic simulator of the glucoregulatory system and propose a particular implementation strategy based on reducing the frequency of the observations and rewards passed to the agent, and using a simple reward function. We train agents with that strategy for three groups of patient classes, evaluate and compare it with alternative control baselines. Our results show that our method is able to outperform baselines as well as similar recent proposals, by achieving longer periods of safe glycemic state and low risk.

Introduction

Diabetes mellitus (DM) is a disease associated with abnormally high levels of blood sugar (blood glucose) due to lack of insulin (type 1 diabetes - T1D) or insulin resistance (type 2 diabetes - T2D). In 2019, approximately 463 million adults worldwide were suffering from DM, which is increasing continuously¹. More than 1.1 million children and adolescents are living with type 1 diabetes. In addition, there are 4.2 million deaths caused by DM.

T1D can produce little or no insulin, so patients must monitor their blood glucose (BG) level and manually administer insulin doses. Excessive insulin doses, may cause *hypoglycemia*, that is, low blood glucose levels, which can lead to short-term complications, such as drowsiness, shakiness, confusion, loss of consciousness or a seizure², or even coma or death³. On the other hand, too little insulin can result in *hyperglycemia*, that is, high blood glucose levels and can cause long-term chronic diseases, including retinopathy, nephropathy, and neuropathy³.

A Continuous Glucose Monitor (CGM) is a sensor, usually a wearable device, that provides real-time readings of BG levels. Combining an insulin pump and CGM with a control algorithm to deliver insulin⁴ is an alternative to the traditional patient self-management of insulin injection to manage BG levels in DM. Control algorithms can be classified into: (1) open-loop controls, which require patient intervention and/or external information, such as meal or exercise announcement; and (2) closed-loop controls, which do not require the patient intervention to regulate the dosage⁵ but some external information can be useful in avoiding rapid BG growth⁶. The combination of a CGM, an insulin pump and a closed-loop control algorithm is usually referred to as an Artificial Pancreas (AP). Since the technology for CGMs and insulin pumps is relatively mature, research for AP has been focused in the last years on suitable closed-loop control algorithms⁵⁻⁹.

Several methods for closed-loop controls can be found in the literature, including proportional integrative derivate (PID)⁷ methods, model predictive control (MPC)^{8,10} and expert-based approaches⁹. In brief, PIDs, one the most used solutions⁶, use previous BG samples as feedback to determine the insulin needed to drive the desired glucose concentration in human blood. Their main disadvantage is a poor adaptation to meal disturbances^{11,12} and inability to individual treatment. MPC requires a model to predict future glucose concentration using current BG, insulin delivery and meal intake; then the algorithm calculates the appropriate insulin infusion rate by minimizing the difference between estimated glucose concentration from the model and the target glucose concentration on a prediction time window⁸. These methods depend obviously on the quality of the

model and are also sensitive to external disruptions such as food intake or physical activity, that cannot be accurately modeled. Expert-based approaches implement case-based logic using experience from medical experts to decide when and how much insulin to deliver⁹. The model requires an expert to create, adjust and evaluate the model, which can lead to human errors. Also, existing empirical models of a patient metabolism cannot be applied to these approaches, hence there are no theoretically based performance guarantees.

Artificial intelligence and machine learning techniques are increasingly being used in medicine and healthcare. Among them, Artificial Neural Networks (ANNs) involve machine learning algorithms that have been already used for diabetes purposes. They need labeled training data from experts to predict blood glucose concentration based on supervised learning and avoid human error. ANNs perform well for short-term prediction¹³. However, improving the prediction of supervised learning approaches implies high volumes of training labeled data and there still remains the problem of designing and appropriate controller from the predicted BG level. Reinforcement Learning (RL) has been suggested as a more promising alternative^{5,11}. In RL, a software agent makes observations and takes actions within an environment and receives rewards from its actions. By appropriately shaping the reward function, the agent can self-learn the desired goal. Their main advantage is that they are model-free, up to some extent since the environment provides the implicit model, and can learn latent disturbances and adapt to them. Nowadays, there are realistic simulators of the glucoregulatory system, even approved by the FDA¹⁴, that can be used as an environment in an RL framework¹⁵.

Even though different RL approaches have been increasingly proposed and discussed^{5,11,12}, effective training of agents for BG control has proved to be difficult^{10,11,16,17}. Several factors may explain the difficulties. First, most of the RL algorithms are designed to approximately solve a Markov Decision Process (MDP) with a fully observable state space¹⁸, but realistic simulators of the glucoregulatory systems cannot be considered fully observable. Therefore, the environment is at most a Partially Observable MDP (POMDP) and, in spite of their ubiquity in many fields, only very few recent algorithms have been developed specifically for POMDPs¹⁹. Moreover, POMDPs require a mapping from true environment states to observable variables that has to be defined by the control algorithm designer⁵, usually with a high degree of arbitrariness. This may have a serious influence on the learning ability of the agent because a bad choice of observable states make state changes and associated rewards not directly related to agent actions. For instance, a delayed response to insulin is a realistic feature of the simulator but it is also related to the choice of POMDP state mapping, because a CGM reading (observation) after an insulin injection (action) does not reflect the actual change of state. Second, compared to other learning environments, there is a remarkable number of design alternatives, whose influence is not clear and usually require careful trial and evaluation. Those involve the choice of the RL algorithm (agent) and its underlying neural network architecture, the tuning of agent hyperparameters, the selection of an appropriate reward function and even the design of the action space that may be adapted to patient specific data^{11,17}. For the sake of conciseness, the design choices related to the aforementioned issues are called implementation strategy in the rest of the paper.

From the above discussion, it is clear that different implementation strategies can lead to effective RL-based closed-loop controls. In this paper, we evaluate two different reinforcement learning algorithms to control *in silico* blood glucose levels in T1D and compare them with other well-known alternatives, including a PID controller. We describe and discuss our implementation strategy and related problems and compare it with recent proposals with a different implementation strategy^{11,17}. Our paper shares with those proposals the initial premise: to train state-of-the-art Deep Reinforcement Learning (DRL) algorithms, such as SAC²⁰ or PPO²¹, as a T1D BG closed-loop control, but differ in several aspects of the implementation strategy that we discuss in the paper. Our results show that our strategy can effectively control BG levels, outperforming control baselines in terms of fraction of time spent in the desired glycaemic state and risk metrics. All the source code used for evaluation and training as well as the trained agent policies described in this paper are publicly available in our repository (<https://github.com/girtel/AIML4Diabetes>).

The paper is organized as follows. Next section contains background information needed to better understand the contribution. Afterwards we describe our implementation strategy and design methodology. Our results are described and compared with baselines in the Results sections, and compared with similar proposals in the Discussion section. Finally, we provide concluding remarks.

Background and Related Work

Type 1 Diabetes Mellitus

Type 1 diabetes (T1D) is an autoimmune disease involving that the pancreatic islets produce little or no insulin. Insulin is an anabolic polypeptide hormone regulating carbohydrate metabolism and normal glucose level in the circulatory system. Pathogenesis of T1D is not yet known²², but it is often found in family members who have had a history of this disease. T1D is characterized by recurrent or persistent excessive amount of glucose in blood plasma, that is, hyperglycemia, which can be diagnosed by several criteria defined by World Health Organization (WHO)²³. Patients with this type of diabetes need to use insulin by external means to maintain blood sugar levels, usually subcutaneous injection. There are four main types of

Figure 1. Principal components of the computer simulation environment.

injectable insulin: rapid-acting insulin, short-acting insulin, intermediate-acting insulin, and long-acting insulin. The insulin dose that a patient has to inject is traditionally planned by a physician according to the patient characteristics and clinical history.

Models and Simulation

Currently, it is possible to experiment with virtual patient populations in simulators to determine the optimal insulin dosage before being used in patients^{14,24}. Several simulators exist to research on T1D models, such as TD1_VPP²⁴, Dosing-RL Gym²⁵ or UVA/PADOVA Simulator¹⁴. TD1_VPP is a virtual patient population software which can generate virtual patients following single-hormone and dual-hormone mathematical models, integrating the effects of exercise in the glucoregulatory system. Dosing-RL Gym is based on an expanded version of the Bergman minimal model, which includes meal figure 1 disturbances²⁵.

The UVA/PADOVA Type 1 Diabetes Simulator can be used as a substitute for preclinical testing of closed-loop control strategies. The simulator is developed by the Universities of Virginia (UVA) and Padova and has been accepted by the US Food and Drug Administration (FDA)¹⁴. It is the most used simulator among *in silico* software according to⁵ and²⁶. There are four main components of the simulation, which are depicted in Fig. 1: (1) *In silico* patient – a model of the glucose-insulin system in a patient; (2) *In silico* sensor – a model of sensor error; the controller to be tested; (3) Controller – a model used to estimate the amount of insulin to maintain blood sugar; and (4) *In silico* pump – a model of discrete insulin delivery and subcutaneous kinetics. In this paper we use an available open-source Python implementation of the UVA/PADOVA simulator¹⁵, previously used in similar studies^{5,11}. The simulation environment implements the OpenAI gym interface²⁷, which makes it compatible with several RL frameworks. It also simulates different noisy CGM sensors, insulin pumps and a random meal scheduler.

Deep Reinforcement Learning

Reinforcement Learning (RL) refers to a class of methods to optimize the decision process of an *agent* operating on a given *environment*. The decision process is usually considered to be a *Markov Decision Process* (MDP), a discrete-time stochastic control process. Formally, an MDP is defined by a 4-tuple (S, A, P, R) of states, $s \in S$, actions, $a \in A$, a state-transition function $Pa(s, s') = Pr(s_{t+1} = s' | s_t = s, a_t = a)$, and reward function $R_a(s, s')$. The agent is the learner entity that seeks the optimal

behavior and is able to perform an action $a(s)$, which changes the state. In this change of state from s to s' , the agent obtains a reward r , considered as the feedback from the environment. MDP-solving algorithms employ what is called a policy, denoted as π , which is a mapping between states and actions; that is $\pi : s \rightarrow a$. Their goal is to reach an optimal policy π^* , which maximizes the accumulated sum of rewards over the entire lifespan of the agent. This decision policy can be determined by the state-action function, also called Q-function, $Q(s, a)$, which can be approximated using Deep Neural Networks (DNN). Deep Reinforcement Learning (DRL) refers to algorithms and methods that use DNN to approximate the Q-function or optimal policy.

It is commonly assumed that the MDP has a fully observable state space S , that is, that the agent has access to observations that fully represent the state of the environment. However, the observation may just be a partial representation of the underlying state. A Partially Observable Markov Decision Process (POMDP) is an extension of an MDP, where the agent cannot fully observe the system state. In that case, the MDP 4-tuple is extended with a space of observations, $o \in O$, and a usually unknown and potentially stochastic function that maps the observations to true underlying states, $T : o \rightarrow s$. Partial observability may stem from many factors, including limited sensing capabilities or unknown environment dynamics¹⁹. Let us remark that partial observability can also be applied to the reward function²⁸. Despite the ubiquity of POMDPs in many practical systems¹⁹, most of the DRL algorithms assume an underlying MDP^{18,21}. POMDPs are usually addressed in DRL by augmenting the observation space with the history of past observations and actions and the use of Recurrent Neural Networks (RNN)¹⁹ in the architecture of the learning algorithm. Only recently a few algorithms have been specifically designed to deal with POMDPs¹⁹. Moreover, the state transition in some environments is determined not only by agent actions, but also by exogenous stochastic input actions²⁹. Efficient methods to deal with this kind of environments are discussed by Mao *et al.*²⁹.

Most of the state-of-art DRL algorithms used for MDPs are based on actor-critic methods: temporal difference learning algorithms that separate representations of value functions and policies explicitly¹⁸. The actor selects actions in the action space, while the critic estimates the value function from the action made by the actor. They can be applied to either discrete or continuous action spaces. However, these methods show poor sample efficiency and stability convergence properties. A variety of techniques have been developed to address those problems^{18,21,30}. In this paper, we have used two of them: Soft Actor-Critic (SAC), a widely used continuous-state DRL algorithm^{11,20}, whose policy maximizes a trade-off between the expected return and entropy, a measure of randomness in the policy, which ensures higher robustness and stability¹⁸. Maximum entropy policies have been shown to solve POMDP with unobserved rewards²⁸. Besides SAC, we have used Proximal Policy Optimization (PPO), another popular DRL algorithm²¹. PPO ensures that its policy does not change much from the previous policy updates, leading to smooth learning and avoiding variance in training. The tradeoff between SAC and PPO is stability and sample efficiency. PPO tends to be more stable and uses more data, whereas SAC tends to be the opposite. PPO is also claimed to work well on POMDPs³¹. Both allow the use of RNNs in their architecture.

Related Work

State-of-the-art control algorithms for AP systems on devices already on the market are based on PID and MPC approaches^{6,11,32,33}. Two commercially available FDA-approved systems are Metronics 670G and 770G using PID³²⁻³⁴ and Tandem Control-IQ using MPC^{32,33}. PID controller are simple but have problems to adapt to meal consumption^{6,11}. MPC controllers are more proactive than PID in insulin delivery by predicting BG levels. PID and MPC controllers usually work as hybrid closed loop system, requiring announcements of meal carbohydrates amount and exercise activity³³. Most commercial products avoid hypoglycemia overnight by utilizing Predictive Low Glucose Suspend (PLGS)³³.

The number of data-driven models for prediction of BG in T1D is increasing²⁶. Reinforcement learning is being used in recent research works in the field of health care. For instance, the RL Q-Learning algorithm was applied on discrete action space simulation for radiotherapy to understand scenarios of tumor growth and its treatment plan³⁵. RL agent-based models have been used on continuous state and action spaces problems to find cytokine therapy for sepsis, reducing mortality from 49% on average to 0.8%³⁶. RL is suitable for time sequence problems as in the glucoregulatory system. Furthermore, the agent can learn the policies without the need for labeled data as in supervised learning methods^{5,12}. The ability of RL to capture food intake patterns without human input, makes it a good candidate for a fully closed-loop system and more responsive and safer policies¹¹.

Regarding specifically T1DM control, a recent review discusses most of the approaches used so far in this topic⁵. That review exposes the wide variety of alternatives used in almost all of the defining elements of the RL framework such as the definition of the state space, the action space, class of RL algorithms used and the reward function, what we have called the implementation strategies. To mention a few which differ in the class of RL algorithms, in¹² a control based on RL is proposed where the value function is not approximated by a DNN but by a quadratic function and in³⁷ the value function is approximated by a Gaussian process. In both cases, robust solutions are provided, but simplified glucose models are used. On the contrary, in this paper we use a more realistic simulator which generally requires DRL to approximate the value function.

Fox *et al.* adopt very similar approaches to the one in this paper in two recent papers^{11,16}. In fact, they share the basic

premise with this paper: training a DRL agent for BG closed-loop control. However, both this work and their previous papers employ different implementation strategies. In their first one¹⁶, the state space consists of the previous 24 hours of CGM samples and insulin doses at 5-minute intervals, but the action space is discrete and made of only 3 insulin doses. Three relatively simple DNN architectures were used to approximate the value function and the improvements over the PID baseline were not particularly noticeable. In their second one¹¹, the state space is also made of CGM and insulin samples but only from the 4 previous hours and the action state is continuous and so the SAC algorithm used. Additionally, the reward functions differ in both cases. In contrast, in this work we use only the last CGM sample as state but take a relatively long interval of 30 to 60 min between observations and actions and define a simpler reward function. A hybrid model-based and DRL approach is discussed by Yamagata *et al.*¹⁰, which uses a discrete action space combined with meal announcement. Finally, very recently Lin *et al.*¹⁷ proposes a combination of machine learning methods for BG control: the controller uses a DRL SAC agent which is driven by a PID control as an initial policy and, in addition, the observation state is extended by the predictions of a dual attention network. Finally, the actions are also regulated by a adaptive safe action.

The results of the last three aforementioned methods are further compared with ours in the Discussion section.

Implementation Strategy and Methodology

In this section we describe and discuss the design process and choices for our implementation strategy for a BG closed-loop control based on DRL. Our goal is to balance blood glucose as longer as possible with low risks. First, we define the state and action space and discuss the reward function. Afterwards, we conduct an initial evaluation based on naive strategies to determine the features of the environment that may have more influence on the agent learning. With these results we refine our design and propose an implementation strategy that is evaluated in the Results section.

Analysis of environment and initial design

The simglucose simulator environment implements the UVA/Padova glucose model³⁸ and provides CGM sensors that sample the BG level through a noisy (stochastic) process as well as a random, patient-dependent, meal schedule. From this description, it is clear that the environment should be considered a POMDP, since the CGM observations of BG levels include noise from sensors. In fact, the underlying hidden states, s , of the environment are given by the glucose model states³⁸, rather than the BG level. Moreover, the dynamics of the environment, that is, the state transitions, are not only determined by the actions taken by the agent (insulin dose injected), but also by exogenous stochastic input actions²⁹, such as the intake of CHO (meals), physical exercise and others. Since simglucose does not consider physical exercise, Unlike T1D_VPP²⁴, we restrict our study to external meals.

All of the above considerations have an influence on the design of the DRL controller. We first define the elements of the DRL algorithm for our problem as follows:

Observations and state: The mapping of the observable variables may determine whether the environment is a POMDP or an MDP. As an example, the number of frames included in the observations in the Atari Pong game makes it become a POMDP or a MDP³⁹. We start by using only CGM samples as observation variables since they are readily available. Unlike¹¹ we do not use past actions (insulin doses) in the observations. Since we target a closed-loop controller, we do not include CHO intake as part of the observation, which has to be announced by the user, even though some devices may facilitate its announcement⁴⁰. We start by using only the current observation, given by the last CGM sample, $o \in \mathbb{R}^+$. The CGM sample frequency is three minutes per environment step.

Action: The action is the amount of basal insulin that the patient gets injected. It is a decimal number ranging between 0 to 30 units, $a \in [0, 30]$, according to the specifications of the insulin pumps implemented in the simulator.

Goals and risk metrics: Safety is crucial in healthcare applications. The main goal of our BG controller is to balance the BG level for as long as possible with low health risks. A commonly used metric of risk associated with BG levels is the blood glucose risk index (BGRI), and it has also been used to measure the performance of control methods⁴¹. BGRI is a measure of glucose variability and associated risks and it is based on a symmetrization of the BG measurement scale⁴¹. The Clarke BGRI is defined as $BGRI = LBGI + HBGI$, where $LBGI$ and $HBGI$ are computed over a series of n CGM samples as:

$$LBGI = \frac{1}{n} \sum_{t=1}^n rl(BG_t) \quad (1)$$

and

$$HBGI = \frac{1}{n} \sum_{t=1}^n rh(BG_t) \quad (2)$$

Figure 2. Blood Glucose Risk Index Function. Our target value is BG level at 112.517 mg/dL for Clarke RI and 138.89 mg/dL for Magni RI and having a zero BGRI implies that there is no risk for a patient at this point (Zero-risk). Figure shows graphically this pint, including the left and right sides of the zero-risk target value, i.e., LGBI and HBGI, respectively. BG level at 70-350 mg/dL is considered as the target range.

where *LGBI* and *HBGI* represent the risk associated with low and high BG levels. They are computed from the following function:

$$f(BG) = 1.509 \times [\log(BG)^{1.084} - 5.381] \quad (3)$$

Noted that BG is measured in mg/dL. $f(BG)$ is the basis to calculate the BG risk function using the formula $r(BG) = 10 \times f(BG)^2$ and separating it as low *rl* and high *rh* as follows:

$$rl(BG) = \begin{cases} r(BG), & \text{if } f(BG) < 0. \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

$$rh(BG) = \begin{cases} r(BG), & \text{if } f(BG) > 0. \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

Similar risk measures have been defined and used in the literature. For instance, in¹¹ the Magni risk function is used, defined as

$$ri(BG) = 10[3.35506((\ln(BG))^{0.8353} - 3.7932)^2] \quad (6)$$

Both curves are shown in Fig. 2. As can be seen, the BGRI adequately captures the increased risk associated with hypoglycemia for the patients.

It is common practice to define the reward function in terms of the Clarke or Magni RI^{11,16}. But, as we discuss next, they may not be adequate to capture our intended goal.

Termination limits and safety: In DRL an environment finishes when some condition is met. For instance, in standard OpenAI gyms such as the BipedalWalker⁴² the episode finishes when the robot falls. In simglucose, the episode ends when the BG level goes out of a predefined range. T1D patients should aim for a target range of 70–180 mg/dL⁴³. We have configured simglucose to end the episode when BG<70 mg/dL or BG>350 mg/dL in order to try to avoid dangerous BG levels. Let us note that this is a quite conservative range. In contrast, in¹¹, episodes are done when the BG falls below 10 mg/dL or raise over 1000 mg/dL.

Reward: A crucial, and problematic, aspect of DRL is the need to design a reward function that helps the agent learn the intended goal³⁹. In our design of the reward function, we have tried different rewards. In particular, we have tried using the negative of the Clarke BGRI, in order to keep the BG at the desired, level but negative rewards in this case induce the agent to terminate early, leading to dangerous regimes. A termination penalty to correct for this behavior is usually introduced, as is done in¹¹. We have tried a different strategy: combined with the conservative termination limits mentioned above, we use a simple reward to encourage large episodes, which results in extended periods within adequate BG level regimes. Therefore, the reward is simply:

$$reward = \begin{cases} 1, & \text{if } BG \in [70, 350] \text{ mg/dL.} \\ -100, & \text{otherwise.} \end{cases} \quad (7)$$

Discount factor: the discount factor helps to balance the importance of immediate and future rewards. Since the effect of insulin on the BG is usually delayed, we set it to a relatively large value of $\gamma = 0.999$.

Initial evaluation

We conducted a series of initial tests to determine the features of the environment that may had a greater influence on the agent learning, according to the initial choices described in the previous section. Simglucose comes with a population of 30 virtual patients: 10 adolescents, 10 adults, and 10 children, which statistically represent different cohorts of patients¹⁵. We have tested on one patient from each group the initial alternatives that are summarized below and in Table 1.

Table 1. Summary of Initial evaluation. All columns show average $\pm$ 95% confidence intervals.

Subject	Alg/Obs/Reward	Length (min)	Hypoglycemic (%)	Hyperglycemic (%)	Euglycemic (%)
Child#1	PPO/O1/R1	234 $\pm$ 50.921	0.0077 $\pm$ 0.003	0.44 $\pm$ 0.03	0.550 $\pm$ 0.040
Child#1	PPO/O2/R1	213.3 $\pm$ 34.836	0.004 $\pm$ 0.004	0.41 $\pm$ 0.06	0.58 $\pm$ 0.069
Child#1	PPO/O1/R2	42.9 $\pm$ 9.8369	0.074 $\pm$ 0.009	0 $\pm$ 0	0.925 $\pm$ 0.0093
Child#1	PPO/O2/R2	41.1 $\pm$ 6.7217	0.075 $\pm$ 0.007	0 $\pm$ 0	0.92 $\pm$ 0.0079
Adolescent#1	PPO/O1/R1	617.7 $\pm$ 155.75	$\pm$ 0	0.48 $\pm$ 0.11	0.510 $\pm$ 0.11
Adolescent#1	PPO/O2/R1	536.7 $\pm$ 147.89	$\pm$ 0	0.58 $\pm$ 0.09	0.415 $\pm$ 0.094
Adolescent#1	PPO/O1/R2	59.4 $\pm$ 1.8709	0.050 $\pm$ 0.001	0 $\pm$ 0	0.94 $\pm$ 0.0015
Adolescent#1	PPO/O2/R2	60 $\pm$ 0	0.05 $\pm$ 0	0 $\pm$ 0	0.95 $\pm$ 0
Adult#1	PPO/O1/R1	555.6 $\pm$ 157.16	0 $\pm$ 0	0.57 $\pm$ 0.14	0.42 $\pm$ 0.1
Adult#1	PPO/O2/R1	505.5 $\pm$ 121.09	0.0057 $\pm$ 0.001	0.41 $\pm$ 0.08	0.57 $\pm$ 0.08
Adult#1	PPO/O1/R2	63.6 $\pm$ 0.85843	0.047 $\pm$ 0.000619	0 $\pm$ 0	0.95 $\pm$ 0.0006
Adult#1	PPO/O2/R2	63.6 $\pm$ 0.85843	0.047 $\pm$ 0.000619	0 $\pm$ 0	0.95 $\pm$ 0.0006
Child#1	SAC/O1/R1	39 $\pm$ 0	0.0769 $\pm$ 9.9276e-18	0 $\pm$ 0	0.92308 $\pm$ 0
Child#1	SAC/O2/R1	63.6 $\pm$ 19.166	0.05248 $\pm$ 0.010184	0.030 $\pm$ 0.064	0.91773 $\pm$ 0.05
Child#1	SAC/O1/R2	36.3 $\pm$ 0.64	0.082 $\pm$ 0.001	0 $\pm$ 0	0.91731 $\pm$ 0.001
Child#1	SAC/O2/R2	36.9 $\pm$ 0.98	0.08 $\pm$ 0.002	0 $\pm$ 0	0.91859 $\pm$ 0.002
Adolescent#1	SAC/O1/R1	123.6 $\pm$ 12.13	0.02 $\pm$ 0.001	0 $\pm$ 0	0.97538 $\pm$ 0.001
Adolescent#1	SAC/O2/R1	97.2 $\pm$ 15.30	0.03 $\pm$ 0.0046488	0 $\pm$ 0	0.96773 $\pm$ 0.004
Adolescent#1	SAC/O1/R2	54 $\pm$ 0	0.05 $\pm$ 4.9638e-18	0 $\pm$ 0	0.94444 $\pm$ 0
Adolescent#1	SAC/O2/R2	64.5 $\pm$ 7.57	0.04 $\pm$ 0.0044488	0 $\pm$ 0	0.95248 $\pm$ 0.004
Adult#1	SAC/O1/R1	87.9 $\pm$ 14.83	0.03 $\pm$ 0.003685	0 $\pm$ 0	0.96466 $\pm$ 0.00
Adult#1	SAC/O2/R1	137.1 $\pm$ 29.87	0.02 $\pm$ 0.005159	0 $\pm$ 0	0.97598 $\pm$ 0.00
Adult#1	SAC/O1/R2	59.1 $\pm$ 0.98	0.05 $\pm$ 0.00086268	0 $\pm$ 0	0.94921 $\pm$ 0.0008
Adult#1	SAC/O2/R2	66.6 $\pm$ 2.68	0.04 $\pm$ 0.0018355	0 $\pm$ 0	0.95481 $\pm$ 0.001

- Algorithms: **PPO** and **SAC**, with a standard configuration using dense DNN with two hidden layers of 256 units. In addition, we used an alternative recurrent architecture intended to capture temporal context, which uses as actor and critic networks a RNN with a 10 LSTM cells. We call this variant **PPO-RNN** and **SAC-RNN**.
- Observation space: we used as observation both (**O1**) the current CGM sample and a (**O2**) vector of the past 20 CGM samples, corresponding to the previous hour at 3-minute intervals.
- Reward functions: we used (**R1**) the one point per step reward of eq.(7) and (**R2**) the negative of the Clarke BGRI with a termination penalty. In all the case we set the termination limits to $BG < 70$ or $BG > 350$.
- Meal schedule: the simulator selects a non-deterministic meal schedule particular to each patient according to the Harris-Benedict algorithm¹¹.

We used the PPO and SAC implementations from *stable-baselines*³⁴⁴ and PPO-RNN, SAC-RNN from *TensorFlow Agents* (TF-Agents⁴⁵) and trained it on two Nvidia GeForce RTX 2080 GPUs and Intel Core i9-10920X CPU @ 3.50GHz 12 cores. All agents were trained for 1 million steps, keeping the best model (best average reward) and with a maximum episode length of 10 000 steps (a step represents 3 minute interval).

From the average length of the episode, it can be seen that the agents are not able to keep a safe control of BG levels. In general, PPO is able to achieve longer duration within a safe BG range, because it injects lower insulin levels and patients spend more time in hyperglycemic state. In fact, most of the episodes with PPO end because of high BG levels. On the contrary, SAC tends to inject insulin very aggressively and patients go very quickly to a hypoglycemic state, ending the episode. In both cases, using just the current CGM observation (O1) or the last one-hour interval CGM observations (O2) has little influence and using the simple reward (R1) works better than the negative of the Clarke BGRI.

Discussion and refinement of design

From the previous evaluation, we hypothesize that the three main characteristics of the simulation that affect the lack of training success are: (1) The action can only drive the BG level to lower values. The rise of BG level depends on the environment dynamics and a consequence of the exogenous input actions, that is, the intake of CHO in the meals. (2) The observation (CGM) is a noisy sample of the BG level. (3) The effect of actions in the BG level is delayed. That is, applying an action does not immediately decrease the BG level. Even though those effects are expected from the initial analysis of the environment and not particularly surprising we think deserve further discussion.

Regarding (1), the agent should learn the policy to deal with this fact, that is, that it does not need to deliver insulin when the BG level is low. In fact, our results show that the PPO agent is able to learn policies that anticipate the meal consumption and the subsequent rise of BG. However, they are not enough precise to control adequately the BG levels. This is probably due to the conservative BG range that leads to early episode ending. Regarding (2), let us just notice that the reward is usually assigned according to the actual BG level, not the CGM sample, which seems to negatively affect learning since it makes termination penalties appear random. But even using only BG levels, instead of noisy CGM samples, did not improve learning. Regarding (3), both the large discount factor and the recurrent architecture should have improved learning. But the termination limits of the environment, the glucose dynamics and the randomness of meal schedules make it hard to learn: if the agent tries high insulin doses, the patient goes very quickly to hypoglycemia and the episode ends, and when the agent injects low doses, meals, which are not included in the observation, raise the BG level ending also the episode. Therefore, it seems that the combination of all these factors prevent the agent from properly learning. Both (2) and (3) stem from the fact that we are dealing with a POMDP and the recurrent architecture should improve learning, but several variations tried in our tests did not actually improve much. We could have tried changing the RNN architecture and other hyperparameters, but due to the large parameter space, we chose to focus on the delay of actions as follows.

Observation frequency and insulin response: We just configured the environment to decrement the frequency of the observations and actions, instead of using the usual 3 or 5-minute CGM sample resolution. Even though the samples are taken, since the simuglucose environment is running and updating the state in *mini-steps* of 1 min, they are passed less frequently to the agent, which then provides an action. The rationale is simple: to let the agent observe the actual effects of its actions, that is, the actual patient insulin response according to the glucose dynamics, instead of a seemingly random transition. Even though it is similar to using as state a history of the previous CGM data, as is done in¹¹, it improves the learning process and leads to more adequate insulin regimes in our results. Therefore, we have introduced an additional hyperparameter, *observation frequency*, which is actually already present in the simulator (CGM sample resolution) although usually left at the default device value (a CGM sample every 3 minutes)¹¹. Let us notice that this new hyperparameter does not increase the complexity compared to using a history of previous samples¹¹, since in that case, the lengths of the history vectors are also additional hyperparameters to be tuned.

Figure 3. Insulin response time for each patient. The insulin dose depicted in the color legend on right was injected at 9:00 AM and no meal is taken subsequently.

We selected the frequency for the observations by testing the insulin response time, when injecting a given unit of insulin without taking any meals. This delay is different depending on the patient group, as expected, and shown in Fig. 3: ten different amounts of insulin doses were used to estimate each subject insulin reaction: from 1 to 30 units. As can be seen, the adult response to insulin tends to be more stable and less pronounced and the reduction in BG starts to show around 45 minutes after injection. Adolescent reaction is slightly slower and more pronounced. Finally, children clearly react faster and more strongly to insulin. In fact, high insulin doses quickly drive some of the patients to hypoglycemia and episode termination. From this observations we chose reduce the frequency of observations as follows: for adults, observations are made every 45 minutes (corresponding to 15 three-minute environment steps), adolescent frequency is set to 1 hour (20 steps) and children is set to 15 minutes (5 steps).

Then, we started over the training of the **SAC**, **SAC-RNN**, **PPO** and **PPO-RNN** agents. After running several experiments only **PPO-RNN** could learn effective policies. Therefore, we have selected this algorithm and architecture for a full evaluation of performance in the Results section.

Summary of implementation strategy

We provide here a brief summary of our implementation strategy, before conducting an evaluation and comparison of alternatives in the next section.

From the discussion in the previous section we derive the following implementation strategy: (1) we reduce the *observation frequency* of the samples of the environment state (CGM) and rewards passed to the agent, depending on the subject (45 minutes, 1 hour and 15 minutes for adults, adolescents and children, respectively), (2) we set broad termination limits for the episodes, $BG \in [10, 1000]$, to let the agent explore more thoroughly the environment; and (3) we use the simple reward function of eq. (8) below, to force the agent to learn to keep the patients in euglycemia for as long as possible.

$$reward = \begin{cases} 1, & \text{if } BG \in [70, 180] \text{ mg/dL.} \\ -100, & \text{otherwise.} \end{cases} \quad (8)$$

Let us note that the reward is accumulated during all the simulation *ministeps* and then passed to the agent. For example, if

Table 2. Hyperparameters for PPO-RNN

Hyperparameter name	Value
actor_fc_layers	200, 100
value_fc_layers	200, 100
actor_lstm_size	128, 128
critic_lstm_size	128, 128
num_environment_steps	25000000
collect_episodes_per_iteration	10
num_parallel_environments	30
replay_buffer_capacity	1001
num_epochs	25
learning_rate	1e-3
num_eval_episodes	20

we set the observation frequency to 10, the environment is going to simulate 10 minimesteps before passing the sample to the agent, and, if BG level has been in the desired range all those minimesteps, the accumulated reward passed will be 10.

Results

Experimental setup

As mentioned in previous sections the UVA/PADOVA simulator provides parameters for fully specifying the glucoregulatory system of patients in three groups: children, adolescents, and adults, each category with 10 patients. Our goal is to keep the patient BG level in the selected range for as long as possible.

We have trained a PPO-RNN agent for each of the patients with the implementation strategy summarized in the previous section and the hyperparameters listed in Table 2. During the training, there are periods of instability, where the average reward drops, as reported in other studies¹¹. We save the policy every 100 training steps and select the policy with a highest average reward as a trained agent. Once trained, the agents are evaluated 20 times with different seeds for all the patients, and statistics for episode length, fraction of time in glycemia states (eu, hypo and hyper) and other metrics are collected. For evaluation, we also set the environment termination limits to $BG = 10$ and $BG = 1000$, in order to test the fraction of time that the patients spend in the different states and make them comparable to similar proposals¹¹. Let us remark that patients whose BG reach levels below or above those limits are considered events that result in serious damage or death.

We have compared our results with two baselines, a basal-bolus regime that simulates the usual self-managed treatment to patients with both T1D and T2D, and a PID baseline.

Basal-bolus Baseline. A multiple daily injection therapy which involves using long-acting insulin with a dosage of $basal = \frac{u2ss \times bodyweight(kg)}{6000(U/min)}$, where $u2ss$ is the steady state insulin rate per kilogram $\frac{pmol}{L \times kg}$, and short or rapid-acting insulin (bolus) to regulate blood glucose concentration, with $bolus = (\frac{carbohydrate}{carbohydrateratio} + \frac{BG_{current} - BG_{target}}{correctionfactor}) / sampletime$. Let us note that this treatment requires the patient to be aware of the meal intake and so it is not a closed-loop control. The controls that require explicit knowledge of meal intake are usually called *controls with meal announcement* in the literature.

PID Baseline. A closed-loop control which uses a discrete PID controller aims to set the system output to a given target, s_t , by setting the control variable a_k as a linear combination of three terms $a_k = K_p P(s_k) + K_i I(s_k) + K_d D(s_k)$, where $P(s_k) = s_k - s_t$, $I(s_k) = \sum_{i=0}^k (s_i - s_t)$ and $D(s_k) = |s_k - s_{k-1}|$.

We use the optimal values for the PID parameters, K_p , K_d , K_i , for each patient provided by Fox *et al.*¹¹.

Episode length

First, we examine the average episode length in evaluation. Table 3 shows the average fraction of episodes that were completed by each patient group. Note that an episode is terminated when the BG level goes out of the 10-1000 mg/dL range, which means that the patient has reached a BG level that may result in serious damage or death. A primary goal, therefore, of any control method is to avoid early episode termination. The PPO-RNN control method is the only one able to make all the patients finish all the 10-day evaluation simulations for all the groups. The adult group is the easier to control and all patients with both PPO-RNN and PID were able to finish the 10-day evaluation period for all the 20 simulations. On the contrary, the children are the most difficult group and BB can only reach on average 75% of full episodes, that is 7.5 days. But being able to finish the evaluation period is not enough to determine the quality of the treatment: BG levels of patients must be kept in the desired range for as long as possible. In the following sections we examine how the controls keep the state of the patient during that period.

Table 3. Fraction of completed 10-day Evaluation Reached for Each Method and group.

Group-Method	Fraction of full episode (average % $\pm$ 0.95 CI)
Adult-BB	97.2 $\pm$ 1.8
Adult-PID	100 $\pm$ 0
Adult-PPO-RNN	100 $\pm$ 0
Children-BB	74.9 $\pm$ 4.8
Children-PID	99.5 $\pm$ 0.3
Children-PPO-RNN	100 $\pm$ 0
Adolescent-BB	88.7 $\pm$ 3.4
Adolescent-PID	100 $\pm$ 0
Adolescent-PPO-RNN	100 $\pm$ 0

Risk Index and Glycemic states

We compare the results of basal-bolus (BB), PID and PPO-RNN controllers for the risk index and fraction of time spent in hyper, hypo and euglycemia. The aim is to determine how well controllers regulate the risk of hypoglycemia and hyperglycemia. First, Fig. 4a shows with boxplots the distribution for all the methods evaluated. As can be seen PPO-RNN makes all patients spend more time in euglycemic state than the baselines, which is the actual goal of this mechanism. Both the median and 25 and 50 percentiles are above those of the other methods. In addition, PPO-RNN also outperforms the baselines globally in terms of fraction of time spent in hyperglycemia and hypoglycemia. It is instructive to remark how the distributions are more informative in this case than single point estimates such as the median or mean. For instance, even though the median for the hypoglycemic fraction are similar for PPO-RNN and PID, we can see that a remarkable number of patients spent an unacceptable large fraction of time in hypoglycemia with PID.

In Fig. 4b we show that PPO-RNN also outperforms the baselines when the patient groups are examined separately. The results for children are especially noteworthy, since as we have discussed, it is the most difficult group to train. For that group we remark that: (1) PPO is able to perform especially well to avoid hypoglycemia, unlike PID, which fails clearly in this aspect; and (2) even though BB seems to provide reasonable results for children, the results of the previous section have to be taken into account, that is, that the control is only able to reach 75% of the full episode on average. In other words, the control may provide a good response on average for typical BG levels and meal intakes, but not be able to react to unusual variations, which drive the patient to a dangerous state.

Figure 4. Comparative of fraction of time spent in different glycemic states

Therefore, to better assess the results, it is necessary to look also at the risk index, which informs us whether the patient is in a safe level within the desirable range. For instance, a patient may spent a large fraction of the time in the euglycemic range but with BG levels very close to the hyper or hypoglycemia thresholds, which may make it vulnerable to unusual conditions, such as irregular meal intakes. Recall that RI penalizes more hypoglycemia, because even though both hypoglycemia or hyperglycemia can lead to fatal outcomes, the short-term effects of hypoglycemia can cause T1D patients to have an immediate crisis³³, as opposed to hyperglycemia, whose effects show in the long term. We show the average RI, HBGI and LBGI all over the evaluation period globally in Fig.5a and by groups in Fig.5b. In both cases, PPO-RNN outperforms the other controls and keeps all the RI metrics within reasonable levels, unlike the baselines, especially PID, which shows high RI metrics. These results show that PPO-RNN keeps the patients within safe limits most of the time, unlike the baselines, which do not success in smoothly control BG levels: Even though the patients may be euglycemic, they exhibit less safe BG levels. This is the reason why, combined with poor adaptability, some BB-controlled patients are not able to finish the 10-day episodes.

Figure 5. Comparative of risk indexes

Discussion

The Results section shows that our implementation strategy can effectively control BG levels and consistently outperforms the baselines. A first result to remark is that our tests showed that no patients controlled by PPO-RNN terminated earlier in any of the evaluation trials. Early termination are called *catastrophic failures* by Fox¹¹, since they signal potentially fatal BG levels. Our work in this aspect shows slightly better results than Fox¹¹, which shows evaluation failure of around 0.1%, but their evaluation extended to 100 10-day replications per patient, whereas ours has been limited to 20 10-day replications. On the other hand, we define a failure when BG goes below 10 mg/dL, whereas they apply an even lower threshold of 5 mg/dL. PPO-RNN clearly improves over Yamagata¹⁰, which reports very high failure percentages, with some patients completely unable to finish any episode. Of course, the critical nature of BG control requires exhaustive evaluation of results. Closed-loop controllers must be evaluated by its capability to keep the BG within an acceptable range for years. As a future work we will evaluate for extended time periods, but, since our tests show that patients keep BG levels above 50 mg/dL and below 400 mg/dL in all the replications so far, we are confident PPO-RNN can avoid failures for longer episodes.

Attending to the time spent in euglycemia, our results are in line with the ones reported by Fox¹¹, but outperforms those of Lim¹⁷ and Yamagata¹⁰. Regarding the former, the implementation strategy, as well as ours, is able to keep euglycemia over 72% of time globally, and, in our case, also for all groups. Our implementation strategy can be considered simpler than the Fox one because our observation space is unidimensional (last CGM sample), which makes training more efficient, whereas Fox uses 96 dimensions (previous 48 CGM and insulin data samples). The extended observation frequency for each patient group

is an additional hyperparameter to be optimized in our case, but the choice of 48 previous samples is also a hyperparameter in their case. Lim and Yamagata papers show only marginal improvements over the baselines using different strategies, and in both cases only report 64% of time in euglycemia. Moreover, in all the discussed papers, the children group shows more difficulties to be appropriately controlled with a closed-loop controller. In fact, in the work by Lim¹⁷ the children group is not even evaluated, despite using a fairly sophisticated control involving DRL driven by PID as initial policy. Similarly, in¹⁰ only 3 patients of each group were evaluated.

The evaluation discussed in this proposal and the aforementioned works show that the application of DRL to BG control is a challenging task. Our results show, however, that DRL has a great potential as a closed-loop controller. In fact, proposals put forward so far discuss just a tiny fraction of the space of potential implementation strategies for DRL-based BG control. The fact that the employed ones at the moment are relatively simple but exhibit a performance better than PID, encourages further exploration of this approach. In our case, there is a great margin for improvement. As a future work, we plan to optimize the hyperparameters of our method and test variations in the RNN architecture. We think that changes in the reward function, narrowing the desired range to obtain reward, may help improve the euglycemia fraction. The introduction of an extended observation frequency in our case has been the key to make agents learn effectively. It is introduced as a new hyperparameter, which may be optimized, but we plan to try strategies to make agents learn it. Once the strategy has been optimized, we plan to conduct a thorough analysis of the way the agent applies insulin doses and compare it with those prescribed by clinical practitioners. The goal is to find whether trained policies employ unusual patterns of dosage that may help clinical practice. Finally, as alternative implementation strategies, it should be worth trying algorithms designed specifically to deal with POMDP and input-dependent environments^{19,29}.

Conclusion

In this work we propose a closed-loop BG level control based on DRL. We discuss the particular characteristics of a realistic simulator of the glucoregulatory system as a training environment for DRL agents and the complexity of their training in this environment. Effective training of such agents can be achieved by very different design choices for the learning process, which we call implementation strategy.

We describe the initial evaluation of several alternatives conducted on a T1DM UVA/PADOVA simulator and, based on the results, propose a particular implementation strategy based on reducing the frequency of the observations and rewards passed to the agent, and using a simple reward function. PPO-RNN agents are trained with that strategy for three groups of patients, evaluated and compared with PID and BB baselines. Our system is able to outperform common PID and BB strategies in overall terms, attending to healthy glycemic states and risk index. A critical discussion of the results and a comparison with several recent works is provided, indicating that our system outperforms current solutions at a lower computational cost. Euglycemia is maintained in 73% of the time, and no early termination events (BG out of range) are reported. Hence, our results show DRL as a promising methodology to implement closed-loop BG control.

There is a still large space of potential strategies and margin for improvement in future research. In particular, as a future work, in addition to variations on our current strategy, we consider the use of novel algorithms for POMDP and exogenous stochastic input actions.

References

1. International diabetes federation diabetes atlas, 9th edition (2019).
2. Bequette, B. W., Cameron, F., Buckingham, B. A., Maahs, D. M. & Lum, J. Overnight hypoglycemia and hyperglycemia mitigation for individuals with type 1 diabetes: How risks can be reduced. *IEEE Control. Syst. Mag.* **38**, 125–134, DOI: [10.1109/MCS.2017.2767119](https://doi.org/10.1109/MCS.2017.2767119) (2018).
3. Khodakaramzadeh, S., Batmani, Y. & Meskin, N. Automatic blood glucose control for type 1 diabetes: A trade-off between postprandial hyperglycemia and hypoglycemia. *Biomed. Signal Process. Control.* **54**, 101603, DOI: [10.1016/j.bspc.2019.101603](https://doi.org/10.1016/j.bspc.2019.101603) (2019).
4. Bekiari, E. *et al.* Artificial pancreas treatment for outpatients with type 1 diabetes: systematic review and meta-analysis. *BMJ (Clinical research ed.)* DOI: [10.17863/CAM.23542](https://doi.org/10.17863/CAM.23542) (2018).
5. Tejedor, M., Woldaregay, A. Z. & Godtliebsen, F. Reinforcement learning application in diabetes blood glucose control: A systematic review. *Artif. Intell. Medicine* **104**, 101836, DOI: [10.1016/j.artmed.2020.101836](https://doi.org/10.1016/j.artmed.2020.101836) (2020).
6. Bothe, M. K. *et al.* The use of reinforcement learning algorithms to meet the challenges of an artificial pancreas. *Expert. Rev. Med. Devices* **10**, 661–673, DOI: [10.1586/17434440.2013.827515](https://doi.org/10.1586/17434440.2013.827515) (2013).

7. Renard, E., Place, J., Cantwell, M., Chevassus, H. & Palerm, C. C. Closed-loop insulin delivery using a subcutaneous glucose sensor and intraperitoneal insulin delivery: feasibility study testing a new model for the artificial pancreas. *Diabetes care* **33**, 121–127 (2010).
8. Magni, L. *et al.* Model predictive control of type 1 diabetes: An in silico trial. *J. Diabetes Sci. Technol.* **1**, 804–812, DOI: [10.1177/193229680700100603](https://doi.org/10.1177/193229680700100603) (2007).
9. Ibbini, M. & Masadeh, M. A fuzzy logic based closed-loop control system for blood glucose level regulation in diabetics. *J. Med. Eng. & Technol.* **29**, 64–69, DOI: [10.1080/03091900410001709088](https://doi.org/10.1080/03091900410001709088) (2005).
10. Yamagata, T. *et al.* Model-based reinforcement learning for type 1 diabetes blood glucose control. In *Singular Problems for Healthcare Workshop at ECAI 2020 ; Conference date: 29-08-2020 Through 08-09-2020* (2020).
11. Fox, I., Lee, J., Pop-Busui, R. & Wiens, J. Deep reinforcement learning for closed-loop blood glucose control. In *Machine Learning for Healthcare Conference*, 508–536 (PMLR, 2020).
12. Ngo, P. D., Wei, S., Holubová, A., Muzik, J. & Godtliebsen, F. Control of blood glucose for type-1 diabetes by using reinforcement learning with feedforward algorithm. *Comput. Math. Methods Medicine* **2018**, 1–8, DOI: [10.1155/2018/4091497](https://doi.org/10.1155/2018/4091497) (2018).
13. Robertson, G., Lehmann, E. D., Sandham, W. & Hamilton, D. Blood glucose prediction using artificial neural networks trained with the AIDA diabetes simulator: A proof-of-concept pilot study. *J. Electr. Comput. Eng.* **2011**, 1–11, DOI: [10.1155/2011/681786](https://doi.org/10.1155/2011/681786) (2011).
14. Visentin, R. *et al.* The UVA/padova type 1 diabetes simulator goes from single meal to single day. *J. Diabetes Sci. Technol.* **12**, 273–281, DOI: [10.1177/1932296818757747](https://doi.org/10.1177/1932296818757747) (2018).
15. Xie, J. Simglucose v0. 2.1 (2018).
16. Fox, I. & Wiens, J. Reinforcement learning for blood glucose control: Challenges and opportunities (2019).
17. Lim, M. H., Lee, W. H., Jeon, B. & Kim, S. A blood glucose control framework based on reinforcement learning with safety and interpretability: In silico validation. *IEEE Access* **9**, 105756–105775, DOI: [10.1109/access.2021.3100007](https://doi.org/10.1109/access.2021.3100007) (2021).
18. Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction* (MIT press, 2018).
19. Meng, L., Gorbet, R. & Kulić, D. Memory-based deep reinforcement learning for pomdp. *arXiv preprint arXiv:2102.12344* (2021).
20. Haarnoja, T., Zhou, A., Abbeel, P. & Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, 1861–1870 (PMLR, 2018).
21. Schulman, J., Wolski, F., Dhariwal, P., Radford, A. & Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
22. Diabetes fact sheet (2020).
23. Diagnosing diabetes.
24. Resalat, N., Youssef, J. E., Tyler, N., Castle, J. & Jacobs, P. G. A statistical virtual patient population for the glucoregulatory system in type 1 diabetes with integrated exercise model. *PLOS ONE* **14**, e0217301, DOI: [10.1371/journal.pone.0217301](https://doi.org/10.1371/journal.pone.0217301) (2019).
25. Salas, N., Ferguson, B. & Zweig, J. Reinforcement learning for personalized medication dosing.
26. Woldaregay, A. Z. *et al.* Data-driven modeling and prediction of blood glucose dynamics: Machine learning applications in type 1 diabetes. *Artif. Intell. Medicine* **98**, 109–134, DOI: [10.1016/j.artmed.2019.07.007](https://doi.org/10.1016/j.artmed.2019.07.007) (2019).
27. Brockman, G. *et al.* Openai gym. *arXiv preprint arXiv:1606.01540* (2016).
28. Eysenbach, B. & Levine, S. If maxent rl is the answer, what is the question? *arXiv preprint arXiv:1910.01913* (2019).
29. Mao, H., Venkatakrishnan, S. B., Schwarzkopf, M. & Alizadeh, M. Variance reduction for reinforcement learning in input-driven environments. *arXiv preprint arXiv:1807.02264* (2018).
30. Fujimoto, S., Hoof, H. & Meger, D. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, 1587–1596 (PMLR, 2018).
31. Zhao, W. *et al.* Research on the multiagent joint proximal policy optimization algorithm controlling cooperative fixed-wing UAV obstacle avoidance. *Sensors* **20**, 4546, DOI: [10.3390/s20164546](https://doi.org/10.3390/s20164546) (2020).

32. Trevitt, S., Simpson, S. & Wood, A. Artificial pancreas device systems for the closed-loop control of type 1 diabetes. *J. Diabetes Sci. Technol.* **10**, 714–723, DOI: [10.1177/1932296815617968](https://doi.org/10.1177/1932296815617968) (2016).
33. Schoelwer, M. J. & DeBoer, M. D. Artificial pancreas technology offers hope for childhood diabetes. *Curr. Nutr. Reports* **10**, 47–57, DOI: [10.1007/s13668-020-00347-9](https://doi.org/10.1007/s13668-020-00347-9) (2021).
34. Nimri, R. *et al.* Feasibility study of a hybrid closed-loop system with automated insulin correction boluses. *Diabetes Technol. & Ther.* **23**, 268–276, DOI: [10.1089/dia.2020.0448](https://doi.org/10.1089/dia.2020.0448) (2021).
35. Jalalimanesh, A., Haghighi, H. S., Ahmadi, A. & Soltani, M. Simulation-based optimization of radiotherapy: Agent-based modeling and reinforcement learning. *Math. Comput. Simul.* **133**, 235–248 (2017).
36. Petersen, B. K. *et al.* Precision medicine as a control problem: Using simulation and deep reinforcement learning to discover adaptive, personalized multi-cytokine therapy for sepsis. *arXiv preprint arXiv:1802.10440* (2018).
37. Paula, M. D., Ávila, L. O. & Martínez, E. C. Controlling blood glucose variability under uncertainty using reinforcement learning and gaussian processes. *Appl. Soft Comput.* **35**, 310–332, DOI: [10.1016/j.asoc.2015.06.041](https://doi.org/10.1016/j.asoc.2015.06.041) (2015).
38. Kovatchev, B. P., Breton, M., Man, C. D. & Cobelli, C. In SilicoPreclinical trials: A proof of concept in closed-loop control of type 1 diabetes. *J. Diabetes Sci. Technol.* **3**, 44–55, DOI: [10.1177/193229680900300106](https://doi.org/10.1177/193229680900300106) (2009).
39. Hausknecht, M. & Stone, P. Deep recurrent q-learning for partially observable mdps. *arXiv preprint arXiv:1507.06527* (2015).
40. Food, U., Administration, D. *et al.* Fda approves automated insulin delivery and monitoring system for use in younger pediatric patients (2018).
41. Clarke, W. & Kovatchev, B. Statistical tools to analyze continuous glucose monitor data. *Diabetes Technol. & Ther.* **11**, S–45–S–54, DOI: [10.1089/dia.2008.0138](https://doi.org/10.1089/dia.2008.0138) (2009).
42. Wei, T. & Webb, B. A bio-inspired reinforcement learning rule to optimise dynamical neural networks for robot control. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, DOI: [10.1109/iro.2018.8594017](https://doi.org/10.1109/iro.2018.8594017) (IEEE, 2018).
43. Battelino, T. *et al.* Clinical targets for continuous glucose monitoring data interpretation: Recommendations from the international consensus on time in range. *Diabetes Care* **42**, 1593–1603, DOI: [10.2337/dci19-0028](https://doi.org/10.2337/dci19-0028) (2019).
44. Raffin, A. *et al.* Stable baselines3 (2019).
45. Guadarrama, S. *et al.* Tf-agents: A library for reinforcement learning in tensorflow (2018).

Acknowledgements

This work was supported in part by Grant TEC2017-84423-C3-2-P (ONOFRE-2) funded by AEI/FEDER/UE [Agencia Estatal de Investigación (AEI), Fondo Europeo de Desarrollo Regional (FEDER), and Unión Europea (UE)] and PID2020-112675RB-C41 (ONOFRE-3) funded by MCIN/AEI/10.13039/501100011033. The work of Phuwadol Viroonluecha was supported by the AEI Formación de Personal Investigador (FPI) Predoctoral Grant PRE2018-084260.

Author contributions statement

P.V. and E.E. conceived the methods, P.V. and E.E. conducted the tests, P.V., E.E. and J.S. analysed the results. P.V. and E.E. wrote the manuscript. All authors reviewed the manuscript.

Additional information

The authors declare no competing interests.