# In Silico Comparison of Continuous Glucose Monitor Failure Mode Strategies for an Artificial Pancreas

**Yunjie (Lisa) Lu** [1 2] and **Abigail Koay**[1] and **Michael Mayo**[1]

**Abstract.** An artificial pancreas is a medical Internet of Things-based system consisting of a continuous glucose monitor, an insulin pump, and a micro-controller. The use of artificial pancreas systems is becoming increasingly popular amongst patients with type 1 diabetes due to its effective ability to allow the patient better control of his/her own blood glucose levels compared to other more standard treatments. In this paper, the problem of missing sensor readings in the glucose monitor data is considered. How should the micro-controller (which adjusts the insulin pump based on monitor readings) react when the glucose monitor stops transmitting for an unpredictable period of time? A strategy that answers this question is called a failure mode strategy. In this paper, several potential failure mode strategies are explored in the context of simulation experiments. Results are presented showing that at least one effective and simple failure mode strategy ($0.5\mu$ & LR$<$72) exists.

## 1 INTRODUCTION

An artificial pancreas (AP) [10] is a real-time, closed-loop insulin delivery system for patients with type 1 diabetes (T1D). It consists of three components: a continuous glucose monitor (CGM) [3], a controller, and an insulin pump. The CGM estimates the wearer's blood glucose level (BGL) by sampling the interstitial fluid in the subcutaneous tissue beneath the skin; a small needle-like sensor that is applied usually to the abdomen or the upper arm is used for this purpose. Readings from the sensor are taken at set intervals, typically every five minutes, and are sent wirelessly to the controller. The controller in turn adjusts the rate of delivery of insulin (or insulin dose size) during the next interval on the basis of CGM readings [8].

In this work, the problem of missing data in the CGM trace is considered [2]. Missing data is a potentially significant problem for an AP because the controller may not have been designed to deal with this situation, and the default failure mode strategy may be to simply turn the pump off or revert back to a preset basal rate [1] [8]. This can lead to risks if the user is unaware of the problem. For example, a reduced insulin dosage at meal times may lead to periods of severe hyperglycemia, increasing the risk of long-term complications, for example, heart attack and kidney damage. On the other hand, if the pump keeps running while the sensor is off, it could potentially overdose the patient on insulin leading severe hypoglycemia such as coma or seizure, which can be more life-threatening than hyperglycemia. Additionally, non-severe hypoglycemia can lead to discomfort, such as anxiety or blurred vision.

Gaps in real CGM traces occur for a variety of reasons. In most cases, they are benign (e.g. the sensor is no longer accurate and needs to be replaced) but there is also the possibility of deliberate malicious interference (for example, [9] illustrate how the wireless connection between CGM and controller can be jammed).

As mentioned, a simple default strategy if connectivity to the sensor is lost is to simply turn the insulin pump off or switch to a preset basal rate. However, this is unsatisfactory for the reasons mentioned above. An alternative idea is to have the controller replace the missing CGM readings with estimates, and then use these for the insulin dose calculations so that the decisions on insulin delivery rate can still be made. Unfortunately, there is currently no effective method for dealing with missing CGM readings in real-time [12], as most effective time series forecasting (or replacement) methods require readings from the future.

Therefore in this paper approaches based on changing the behaviour of the insulin pump when the sensor is down are considered. It is shown via a set of experiments involving virtual patients and a state-of-the-art AP controller that viable alternative failure mode strategies that can replace the simplistic "pump off" strategy exist. Note that imputation of missing glucose values is not a focus here; instead we are concerning with how the pump should behave when the sensor is unavailable.

## 2 BACKGROUND

To perform the comparison of failure mode strategies, a virtual patient simulator along with a modern AP controller based on fuzzy logic was used.

### 2.1 Virtual Patient Simulator

The simulator utilised in this work is an open-source T1D patient simulator Simglucose [11]. The simulator is a re-implementation of the 2008 version of UVA/Padova T1D simulator described by Dalla Man et al. [5] which has been approved by the Food and Drug Administration (FDA) for pre-clinical trials of specific insulin treatments such as control algorithms for an AP. In brief, the simulator models all components of a patient and AP system. There are several choices for CGM type, each of which come with appropriate sensor error models and reading limits. Additionally, several different types of simulated insulin pump are available.

The virtual patient model used by the simulator describes the glucose/insulin subsystem and is modelled using compartments. Important processes in the virtual patient simulator include a model of the gastrointestinal tract, from which the rate of appearance of glucose in the glucose subsystem is determined during simulated meals; the insulin subsystem, which models both the rate of appearance of insulin, its rate of degradation, and its interactions with the liver and

[1] Department of Computer Science, University of Waikato, New Zealand
[2] contact author email: yl606@students.waikato.ac.nz

body tissues; and the production and storage of glucose in the liver and its utilisation in muscle and adipose tissue. Each virtual patient in the simulator is generated randomly from a joint probability distribution over tens of physiological parameters, and ten different adult patients are currently available in the simulator.

The simulator models meals as well, and the timing and size of meals in terms of carbohydrate (CHO) count are also randomly generated from an appropriate probability distribution.

Figure 1 shows a trace of the simulator's output for one virtual patient. The figure shows CGM readings, actual BGL (not always identical due to sensor lag and error), and the insulin pump rate. The normoglycemia range (see Table 1) is also shown.

## 2.2 Fuzzy Logic Controller

The controller component of an AP serves the function of monitoring (and aggregating) incoming CGM sensor trace data, and then using this information to make making insulin pump rate adjustment decisions in real time (for example, to reduce the risk of hypoglycemia induced by a rapid decline in BGL).

Many algorithms have been proposed for solving this control problem, and they can be broadly grouped into control theory-based approaches such as proportional-integral-derivative (PID) controllers; logic-based approaches; adaptive statistical based-approaches (e.g. based on moving averages); and machine learning-based approaches [10]. A drawback of several of these approaches is the need to either set parameters which are patient-specific (e.g. a PID controller has at least three important constants to set) or to run the controller for a period of time in order to train a predictive model.

In order to circumvent these issues (mostly), we implemented the fuzzy logic controller (FLC) proposed by Mauseth et al. [6] and recently updated [7]. The FLC encapsulates expert knowledge about insulin dosing in the form of fuzzy logic rules. The idea is that the fuzzy rules comprise knowledge about what an expert diabetes clinician would do in a given situation where an insulin dose needs to be set. Since it is a knowledge-based approach, the controller is not overly dependent on setting constants or training data. Furthermore, the inputs to the controller are straightforward: the current BGL is the first input; the BGL velocity (change since the last reading) is the second input; and the BGL acceleration (how the velocity is changing) is the third input.

These three inputs are then "fuzzified" (mapped to fuzzy sets), and following that a fuzzy lookup table containing the expert knowledge is consulted for optimal dose calculation. In general, the table is designed such that higher accelerations and velocities lead to higher doses, while low absolute BGL values and strongly negative velocities switch the insulin pump off.

The FLC does have a single patient-specific parameter, the "personalisation factor" (PF), which can range from zero to ten. The PF dictates the aggressiveness of the insulin treatment, with a value of ten corresponding to a very weak treatment (all doses being multiplied by a factor of 0.002), while a PF value of zero is the strongest treatment (all doses scaled by 1.74). Most patients should be expected to have a PF of around five, indicating a scaling factor of 0.92 of the dose computed by the FLC.

# 3 METHOD

In this section, the experimental setup, in particular the methods by which the PF values are set on a patient-wise basis, and how artificial missing data gaps were generated in the simulated CGM trace are described.

## 3.1 Optimised PF value

Each individual has significantly different insulin sensitivity. Some are more sensitive to insulin, so their blood glucose levels drop rapidly in response to the same insulin dose than those who are less sensitive. Thus, the FLC's PF value is necessarily unique for each individual.

To determine the optimal PF value for each adult virtual patient in the T1D simulator, the following procedure was followed: for each possible PF value, the simulator was run for ten virtual days at a sampling interval of five minutes. The amount of time spent in normoglycemia for each patient and each PF value respectively was recorded. The PF value for each adult virtual patient which maximised the time spent in normoglycemia was then selected.

## 3.2 Missing Data Generation

One guideline concerning proper usage of CGM [4] state that about 70-80% coverage of sensor reading coverage is required over two weeks for calculating metrics. Assuming this is the usual case, then 336 hours of sensor use over two weeks can be expected, leaving 68 hours of total CGM sensor gap over two weeks, or 4.8 hours per day on average. Additionally, gaps are usually not scattered random individual sensor readings, but are likely to occur contiguously in blocks since events such as faults tend to last for significant periods of time. On the basis of this assumption, we implemented a method to generate missing data gaps in the CGM trace. Essentially, the program generates a random start time (between midnight to 1900 hours, including mealtime) for each gap as well as a random length for the gap (between 0.5 hours and five hours). See Figure 1 for an example of the trace with gaps.

## 3.3 Failure Mode Strategies

In this section, five different failure model strategies are described. As a baseline, results for an "ideal" set of simulations in which there are no sensor trace gaps are also included. The failure mode strategies are as follows:

1. Pump off
   This strategy simply switches the insulin pump off when sensor gaps are detected.
2. $\mu$
   $\mu$ is defined as the average value of the total insulin dose from the previous one hour. Since there are twelve readings per hour in an AP with a five minute reading interval, the average dose is defined as
   $$\mu = \frac{1}{12} \sum_{n=1}^{12} d_i \tag{1}$$
   where $d_i$ is the insulin dose $i$ adjustments previously. The $\mu$ strategy operates the pump at a constant rate of $\mu$.
3. $0.5\mu$
   On the basis that $\mu$ may produce doses that are sometimes too high, the $0.5\mu$ strategy is calculated in the same way, except that the constant dose size is halved.

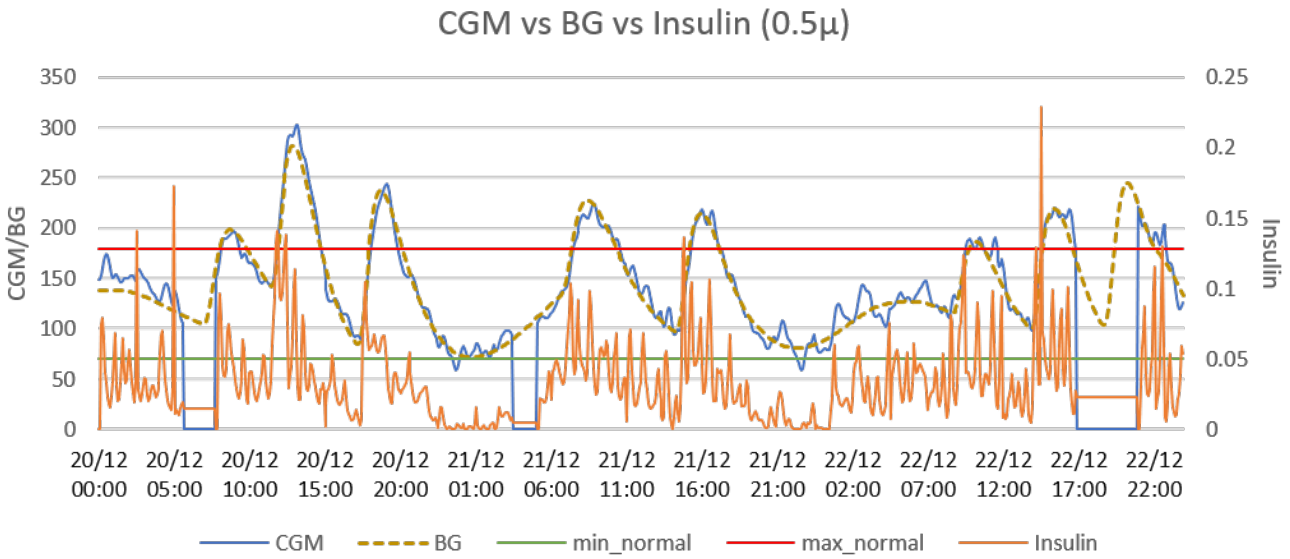


**Figure 1**: Example of data generated by running Simglucose for one virtual patient for 72 hours by using 0.5$\mu$ method. Gaps are shown as periods of zeros in the CGM trace. Not shown here is the CHO intake that primarily drives the fluctuations in BGL.

4. Random choice (Rand_choice)
   Random choice constructs a method that randomly chooses a value from the last one hour of insulin dosages, on the basis that varying the dose size might be beneficial compared to the $\mu$ strategies which computes constant doses. The expected value of the random choice strategy equivalent to the dose as calculated by $\mu$.
5. Hybrid method (0.5$\mu$ & LR<72)
   Following initial results showing that the 0.5$\mu$ strategy was effective, a new strategy was defined that extended 0.5$\mu$ with a constraint: if the last reading (LR) before the sensor is down is smaller than a preset value (in this case, 72.0 mg/dL), then the pump will be turned off; otherwise, the 0.5$\mu$ strategy will be applied. 72 mg/dL was chosen as a threshold because it is just above the L1 hypoglycemia threshold defined in Table 1, and it is desirable that this threshold is not crossed. The equation for this strategy is:

$$\text{dose} = \begin{cases} 0 & \text{if LR} < 72.0 \\ 0.5\mu & \text{otherwise} \end{cases} \quad (2)$$

## 4 EVALUATION

In each experiments, 100 simulated days per adult virtual patient per failure mode strategy was run. This gave a total of 10×6 runs per patient. Each run took approximately ten minutes on a laptop with a 2.6GHz Intel Core i7-9750 CPU processor. For all experiments, the simulated GuardianRT CGM was used in conjunction with the Insulet insulin pump.

Table 2 shows for each patient with their personalised PF value according to the method of selecting the optimal PF value described in the previous section.

The effectiveness of each failure mode strategy was then determined by calculating the time spent in standard glycemic ranges as defined in Table 1. As mentioned, hypoglycemia can be considerably more dangerous than hyperglycemia, so therefore time spent in the L2 Hypoglycemia range is the most significant statistic in the results.

Figure 2 shows that these five strategies provide similar performance in the L1 Hyperglycemia range. All the median values are

| Range | Definition |
|---|---|
| L2 Hypoglycemia | <54 mg/dL (<3.0 mmol/L) |
| L1 Hypoglycemia | 54 to <70 mg/dL (3.0 to <3.9 mmol/L) |
| Normoglycemia | 70 to 180 mg/dL (3.9 to 10.0 mmol/L) |
| L1 Hyperglycemia | >180 to 250 mg/dL (>10.0 to 13.9 mmol/L) |
| L2 Hyperglycemia | >250 mg/dL (>13.9 mmol/L) |

**Table 1**: Definitions of glycemic ranges used by [4].

| Adult# | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|---|---|---|---|
| PF | 6 | 4 | 8 | 6 | 5 | 6 | 7 | 6 | 7 | 6 |

**Table 2**: The optimised PF value for each adult virtual patient.

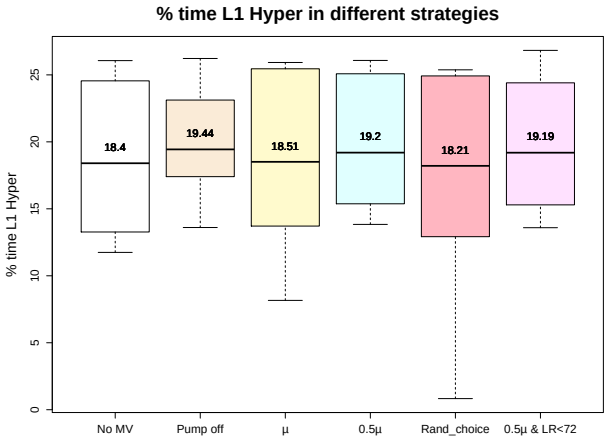


**Figure 2**: Percentage of time spent in L1 hyperglycemia by failure mode strategy across ten adult virtual patients.

quite close to each other (i.e. between 18.21 and 19.44). $\mu$ and Rand_choice perform quite well most of the time. The lowest percentage obtained due to Rand_choice is below 5% for one patient. However, 0.5$\mu$ and hybrid strategy have a more stable performance

with no significant outliers.

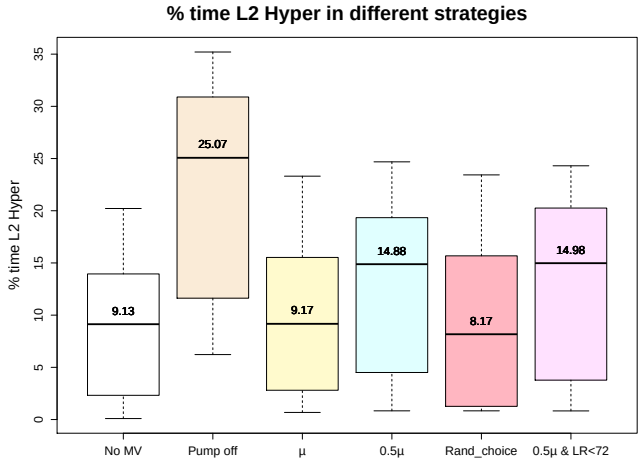


**Figure 3**: Percentage of time spent in L2 hyperglycemia by failure mode strategy across ten adult virtual patients.

Figure 3 depicts time spend in L2 hyperglycemia by the virtual patients. It can be noticed that because of the high chance of being severe hyperglycemia, the pump off strategy is probably the least ideal failure mode strategy. Two strategies that produce the lowest chance of being L2 Hyperglycemic are $\mu$ and Rand_choice.

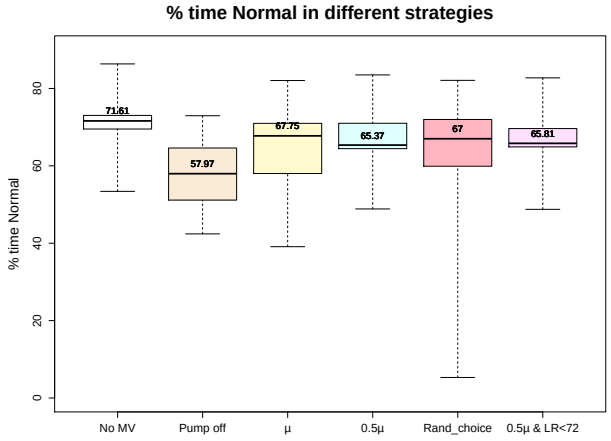


**Figure 4**: Percentage of time spent in normoglycemia by failure mode strategy across ten adult virtual patients.

Figure 4 depicts time spent in normoglycemia. Unlike the other figures where lower is better, in this case higher is better. The figure indicates that the last four strategies output similar median values which are all quite close to the baseline no missing value strategy. The best two strategies that give us the highest median time are again $\mu$ and Rand_choice. However, it is quite noticeable that these two methods, especially Rand_choice, as not stable, as observed before. The percentage time spent in normoglycemia can drop to under 10% for one of these strategies. The $0.5\mu$ & LR$<$72 strategy gives the most stable performance across virtual patients.

Figure 5 and 6 show time spent in the hypoglycemic ranges. From these box plots, it can be observed that both $\mu$ and Rand_choice significantly underperform compared to the other strategies. In terms of the more dangerous L2 hypoglycemia range (Figure 6), the pump off and hybrid strategies work best.

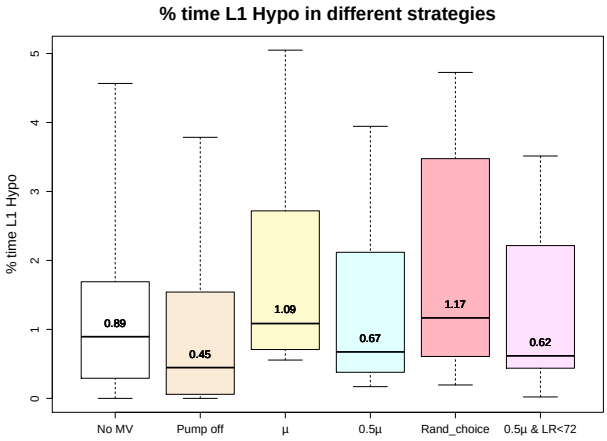


**Figure 5**: Percentage of time spent in L1 hypoglycemia by failure mode strategy across ten adult virtual patients.

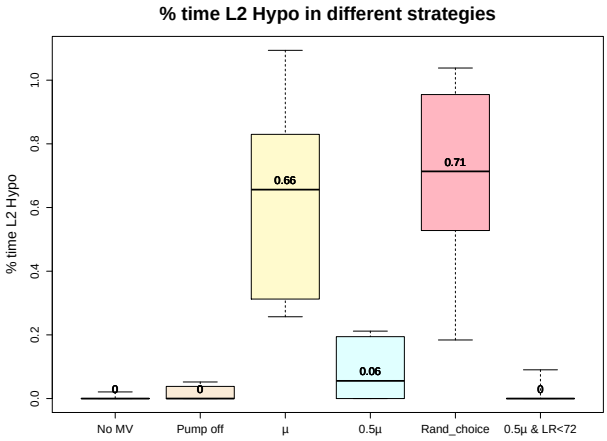


**Figure 6**: Percentage of time spent in L2 hypoglycemia by failure mode strategy across ten adult virtual patients.

## 5 CONCLUSION & FUTURE WORK

In this study, it was found that the hybrid method works well and is furthermore simple to implement. In most cases, it was close to the no missing values baseline strategy, especially in terms of time spent in the critical L2 hypoglycemia range.

The hybrid method could therefore be used as a baseline for the development of more sophisticated failure mode strategies (such as those based on online CGM imputation or machine learning methods) in the future. It is also important to test this failure mode strategy with other controllers and on other twenty patients (i.e., ten children and ten adolescent patients) in addition to the FLC and ten adult patients analysed in this paper.

## REFERENCES


[1] B. Wayne Bequette, 'Fault detection and safety in closed-loop artificial pancreas systems', *Journal of Diabetes Science and Technology*, **8**(6), 1204–1214, (2014). PMID: 25049365.

[2] Helga Blauw, Patrick Keith-Hynes, Robin Koops, and J. Hans DeVires, 'A review of safety and design requirements of the artificial pancreas', *Annals of biomedical engineering*, **44**(11), 3158–3172, (2016).

[3] Giacomo Cappon, Giada Acciaroli, Martina Vettoretti, Andrea Facchinetti, and Giovanni Sparacino, 'Wearable continuous glucose monitoring sensors: A revolution in diabetes treatment', *Electronics*, **6**(3), 65, (2017).

[4] Thomas Danne, Revital Nimri, Tadej Battelino, Richard M. Bergenstal, Kelly L. Close, J. Hans DeVries, Satish Garg, Lutz Heinemann, Irl Hirsch, Stephanie A. Amiel, Roy Beck, Emanuele Bosi, Bruce Buckingham, Claudio Cobelli, Eyal Dassau, Francis J. Doyle, Simon Heller, Roman Hovorka, Weiping Jia, Tim Jones, Olga Kordonouri, Boris Kovatchev, Aaron Kowalski, Lori Laffel, David Maahs, Helen R. Murphy, Kirsten Nørgaard, Christopher G. Parkin, Eric Renard, Banshi Saboo, Mauro Scharf, William V. Tamborlane, Stuart A. Weinzimer, and Moshe Phillip, 'International consensus on use of continuous glucose monitoring', *Diabetes Care*, **40**(12), 1631–1640, (2017).

[5] Chiara Dalla Man, Francesco Micheletto, Dayu Lv, Marc Breton, Boris Kovatchev, and Claudio Cobelli, 'The UVA/Padova type 1 diabetes simulator: New features', *Journal of Diabetes Science and Technology*, **8**(1), 26–34, (2014). PMID: 24876534.

[6] Richard Mauseth, Irl B. Hirsch, Jennifer Bollyky, Robert Kircher, Don Matheson, Srinath Sanda, and Carla Greenbaum, 'Use of a "fuzzy logic" controller in a closed-loop artificial pancreas', *Diabetes Technology & Therapeutics*, **15**(8), 628–633, (2013). PMID: 23829285.

[7] Richard Mauseth, Sandra M. Lord, Irl B. Hirsch, Robert C. Kircher, Don P. Matheson, and Carla J. Greenbaum, 'Stress testing of an artificial pancreas system with pizza and exercise leads to improvements in the system's fuzzy logic controller', *Journal of Diabetes Science and Technology*, **9**(6), 1253–1259, (2015). PMID: 26370244.

[8] Charrise M. Ramkissoon, B. Wayne Bequette Brian Aufderheide, and Josep Vehi, 'A review of safety and hazards associated with the artificial pancreas', *IEEE reviews in biomedical engineering*, **10**, 44–62, (2017).

[9] Luca Reverberi and David Oswald, 'Breaking (and fixing) a widely used continuous glucose monitoring system', in *11th USENIX Workshop on Offensive Technologies (WOOT 17)*, Vancouver, BC, (August 2017). USENIX Association.

[10] Dawei Shi, Sunil Deshpande, Eyal Dassau, and Francis J. Doyle, 'Chapter 1 - Feedback control algorithms for automated glucose management in t1dm: the state of the art', in *The Artificial Pancreas*, eds., Ricardo S. Sánchez-Peña and Daniel R. Cherñavvsky, 1 – 27, Academic Press, (2019).

[11] Jinyu Xie. Simglucose v0.2.1, https://github.com/jxx123/simglucose, last accessed 28 Feb 2020, 2018.

[12] Sara Zulj, Paulo Carvalho, Rogerio Ribeiro, and Ratko Magjarevic, 'Handling missing data in CGM records', in *Future Trends in Biomedical and Health Informatics and Cybersecurity in Medical Devices*, eds., Kang-Ping Lin, Ratko Magjarevic, and Paulo de Carvalho, pp. 420–427, Cham, (2020). Springer International Publishing.