

Metaheuristic Optimization of Insulin Infusion Protocols Using Historical Data with Validation Using a Patient Simulator

Hongyu Wang^{*,§}, Lynne Chepulis[†], Ryan G. Paul^{†,‡} and Michael Mayo^{*}

^{}Department of Computer Science
University of Waikato, Hamilton, New Zealand*

*[†]Waikato Medical Research Center
University of Waikato, Hamilton, New Zealand*

*[‡]Waikato Regional Diabetes Service
Waikato District Health Board, Hamilton, New Zealand
shw168@students.waikato.ac.nz*

Received 29 June 2020

Accepted 13 September 2020

Published 5 November 2020

Metaheuristic search algorithms are used to develop new protocols for optimal intravenous insulin infusion rate recommendations in scenarios involving hospital in-patients with Type 1 Diabetes. Two metaheuristic search algorithms are used, namely, Particle Swarm Optimization and Covariance Matrix Adaption Evolution Strategy. The Glucose Regulation for Intensive Care Patients (GRIP) serves as the starting point of the optimization process. We base our experiments on a methodology in the literature to evaluate the favorability of insulin protocols, with a dataset of blood glucose level/insulin infusion rate time series records from 16 patients obtained from the Waikato District Health Board. New and significantly better insulin infusion strategies than GRIP are discovered from the data through metaheuristic search. The newly discovered strategies are further validated and show good performance against various competitive benchmarks using a virtual patient simulator.

Keywords: Metaheuristic search; continuous optimisation; insulin infusion protocol; simulated validation.

1. Introduction

Diabetes is a fast-growing global problem.¹ It is caused by the body's inability to produce insulin or due to body's resistance to insulin. Insulin is a hormone produced by the pancreatic islets and it promotes the absorption of glucose from blood into

[§]Corresponding author.

This is an Open Access article published by World Scientific Publishing Company. It is distributed under the terms of the Creative Commons Attribution 4.0 (CC BY) License which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

cells. Patients admitted to our local hospital due to diabetes are subjected to continuous intravenous insulin infusion therapy. The goal of such therapy is to keep the blood glucose levels of the patient in a defined normal range (usually 6–10 mmol/L). This goal is achieved by adjusting the insulin infusion rate at approximately hourly intervals in order to administer an appropriate amount of insulin throughout the treatment. Each time, the infusion rate is either adjusted upwards to decrease the risk of *hyperglycemia*, defined as glucose levels being too high, or downwards to minimize the risk of *hypoglycemia*, defined as glucose levels being too low. Infusion rate adjustments are generally carried out by nurses, who refer to a paper-based protocol which includes sets of instructions and a lookup table to determine the appropriate insulin infusion rate measured in units/h (U/h), where one unit of insulin is defined as the “biological equivalent” of 34.7 μg pure crystalline insulin. The protocol also identifies scenarios where a doctor needs to be called to take over immediate care of the patient (e.g. when the patient suffers from severe hypoglycemia).

The research question that our research aims to explore is how to devise new protocols for calculating the optimal insulin infusion rate for these patients in a mathematically consistent way, and how to evaluate such protocols. There are various ways of approaching the problem. In our research, metaheuristic search algorithms are used in conjunction with insulin therapy data obtained from the local hospital. The data come in the form of a time series of blood glucose levels/insulin infusion rates for each patient’s therapy session. The idea is to discover a protocol that *would have* performed better than the current protocol with respect to the historical data, if the protocol being tested had been applied in the same situation.

The Glucose Regulation for Intensive Care Patients (GRIP)² serves as the base framework for the metaheuristic search. The objective of the metaheuristic search is augmented with a methodology in the literature for evaluating new insulin infusion protocols based on historical data.³ GRIP utilizes a simple linear formula that can be applied to data to calculate insulin infusion rate adjustments at any point of an insulin therapy session whenever such adjustments are needed. GRIP requires as its input the blood glucose and insulin infusion data during the past four hours of the therapy.

The resulting protocols obtained from the metaheuristic search are validated using Simglucose,⁴ which is a reimplement of the FDA-approved Type 1 UVA/Padova Simulator.⁴ Simglucose simulates a model of glucose and insulin which represents a patient with Type 1 Diabetes. We use an adapted version of Simglucose to evaluate intravenous insulin protocols in our research.

The novel contribution of our research is twofold. First, it shows that metaheuristic search algorithms can be used to optimize the parameters of GRIP, thus improving and adopting the strategy with respect to the historical data. Second, newly discovered strategies can be validated using simulated virtual patients, who have never been exposed to the optimization process, and exhibit better performance than both GRIP, which is the base of the optimization, and the current local hospital practice, which is the historical protocol used to produce the data for optimization.

This paper extends the paper entitled “Metaheuristics for Discovering Favorable Continuous Intravenous Insulin Rate Protocols from Historical Patient Data”.⁵ The novel contribution of this paper over Ref. 5 is that Simglucose is used to verify optimization results by applying the resulting protocols to virtual patients. This paper shows how to use a patient simulator to fairly evaluate protocols optimized using real-world patient data, which Ref. 5 does not. Additionally, this paper includes more detail than Ref. 5, e.g. more detailed illustration of optimization results and more comprehensive background on insulin infusion protocols.

2. Background

This section covers the background knowledge of our research, including various insulin infusion protocols, metaheuristic search algorithms and Simglucose.

2.1. Existing insulin infusion protocols

In the literature, there are both algorithms developed for in-patients and ones developed for patients in Intensive Care Units. Although our research aims to develop and optimize algorithm-based insulin infusion protocols for in-patients instead of ICU patients, the ICU-based insulin infusion protocols provide nice comparison and some of them can act as starting points for AI optimization.

Table 1 compares the protocols investigated during our research, in terms of the amount of historical data used, the core form of a protocol, and the inner hidden state of a protocol. The protocols are discussed in detail in the following subsections.

It should be noted that the units for blood glucose levels may differ in different protocols, as some protocols use mmol/L while the others use mg/dL. This is heavily dependent on the protocol’s original country, because countries such as the United States use mg/dL for blood glucose measurements, while the rest of the countries use mmol/L, including New Zealand. The units of the protocols are kept as they are in the original publications, and conversion is relatively easy as 1 mmol/L of blood glucose is equivalent to 18 mg/dL of blood glucose.

2.1.1. Local DHB practice

At the local hospital, the nurses use a tabular and textual protocol (shown in Fig. 1 to assist in making intravenous insulin infusion rate decisions, and the protocol’s

Table 1. Comparison of the protocols studied.

	Max history utilized	Core form	Inner state
Waikato DHB	1 h	Instructions	Infusion scale
GRIP	4 h	Formula	None
Glucommander	1 h	Formula	Multiplier
Glucosafe	Cumulative from Start	Formula	Model parameters
Sigmoidal	8 h	Instructions	Infusion scale

Medical Staff –

- to initiate this Protocol on suitable patients who fulfil the criteria
- to determine initial scale (see below) and chart “ IV Insulin Infusion as per Protocol -Scale” onto patient drug chart
- to prescribe 1 unit of NOVORAPID per 10g CHO with each meal
- to prescribe appropriate IV fluids as described on page 1

Nursing Staff –

- to prepare the above procedure
- to adjust the rate and scale of infusion as described below
- to ensure bolus Novorapid is prescribed and given to patients if they are allow to eat

Target CBG - between 4.1 and 8.0mmol/L

6-10 mmol/L

CBG mmol/l	Scale A	Scale B - start	Scale C	Scale D	Scale E
< 4mmol/l	Stop/consult*	Stop/consult* Hold insulin infusion for at least 30 min, when CBG>4.0, resume scale to the left at rate according to CBG - refer to HYPO PROTOCOL			
4.1 – 5mmol/l	0.2 units/hr	0.5 units/hr	1 units/hr	1.5 units/hr	2 units/hr
5.1 – 6mmol/l	0.5 units/hr	1 units/hr	2 units/hr	3 units/hr	4 units/hr
6.1 – 8mmol/l	1 units/hr	2 units/hr	3 units/hr	4 units/hr	5 units/hr
8.1 – 10mmol/l	2 units/hr	3 units/hr	5 units/hr	6 units/hr	8 units/hr
10.1- 12mmol/l	3 units/hr	4 units/hr	7 units/hr	8 units/hr	10 units/hr
12.1-14mmol/l	4 units/hr	6 units/hr	9 units/hr	10 units/hr	13 units/hr
14.1-16mmol/l	5 units/hr	8 units/hr	11 units/hr	12 units/hr	16 units/hr
> 16mmol/l	6 units/hr	10 units/hr	13 units/hr	15 units/hr	20 units/hr

No Change in Scale When:	Move One Scale to the Right When:	Move One Scale to the Left When:
CBG within Target Range	CBG remains above target range (>8.0) and decreasing less than 1mmol/hr	CBG is decreasing MORE than 4mmol per hour
OR	OR	OR
CBG decreasing steadily at a rate of 1-4mmol/hour towards target range	CBG increasing	Hypo occurs (CBG 4.0 or less) see HYPO PROTOCOL

Initial Scale Recommendation:

- **Scale A:** thin, small individuals, especially women; some frail elderly; some very sensitive patients with Type 1 DM (i.e. Total daily dose < 30units/day); athletes; hypopituitarism, hypoadrenal or hypothyroid patients; post-pancreatectomy patients; some well controlled non-obese diet controlled patients with type 2 DM.
- **Scale B: Default scale. Individuals (type 1 or type 2) with no special circumstances or if normal insulin requirements approximately 30-80units/day**
- **Scale C:** Seriously ill (higher fever etc), moderate steroid doses (> 20mg daily of prednisone), very obese or insulin resistant patients (total daily dose > 80units/day)
- **Scale D and E:** upper range and insulin resistance scales. To be used after failure of scale C

Dose and rate of insulin infusion should be recorded on “Intravenous Insulin Infusion Recording Chart” (attached)

2. Waikato Diabetes Service Hospital Guidelines. Dr Peter Dunn

Fig. 1. The insulin infusion rate recommendation protocol currently in use at the Waikato District Health Board, New Zealand. The protocol is in the form of an insulin infusion rate table and text instructions.

table is visualized in Fig. 2. The protocol should be referred to approximately each hour, which means that insulin infusion rate decisions are made hourly.

The table in the protocol represents a function which takes a patient’s blood glucose level and an infusion rate scale as arguments and returns its recommended insulin infusion rate. The scales represent different amounts of insulin recommended

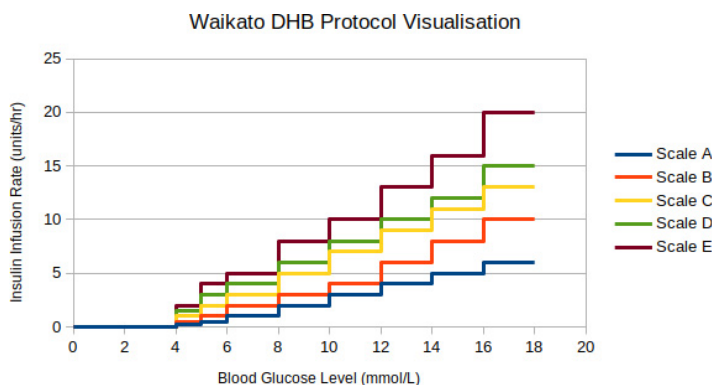


Fig. 2. Visualization of the Waikato DHB protocol.

at each same blood glucose range, with the lower scales such as *Scale A* recommending less insulin while the higher scales such as *Scale E* recommending more insulin. The scales are needed as different patients have different physical conditions, which make them require different amounts of insulin when their blood glucose levels fall in the the same range. For example, a lower-weighted patient generally requires less insulin than a higher-weighted patient when they have the same blood glucose level. In fact, there are many factors that influence a patient's requirement of insulin, and most of them are not as easily identified or quantified as body weight, e.g. liver function, kidney function, insulin sensitivity.

The textual instructions help determine and adjust the insulin infusion rate scale for a patient. The initial scale is determined by a patient's physical condition and medical history, with *Scale B* as the default scale for most patients, *Scale A* for thin and/or insulin sensitive individuals and *Scale C* for insulin insensitive patients, while *Scale D* and *Scale E* are generally not used as the initial scale and are only used when *Scale C* proves to be insufficient for a patient, i.e. when *Scale C* cannot keep the patient's blood glucose level from rising above the normal range and therefore fails to minimize the hyperglycemia risk.

An adjustment of infusion scale may occur if the current scale appear to be either insufficient or excessive for a patient. The ideal blood glucose range is specified according to the patient's condition by the medical staff, and this ideal range plays an important role in determining the insufficiency and excessiveness of an infusion scale. In most scenarios for in-patients at the local hospital, an ideal range of 6–10 mmol/L is selected, which is higher than the 4–8 mmol/L range more commonly seen in ICU settings, largely because the risk of hypoglycemia is usually more significant than the risk of hyperglycemia for in-patients who are already receiving insulin intravenously.

A higher scale is adopted in the event of insufficient insulin, with the insufficiency signified by blood glucose levels above the ideal range and not decreasing at a satisfactory rate, i.e. decreasing by less than 1 mmol/L since last hour or increasing.

On the other hand, a lower scale is adopted when insulin appears excessive, which is signified by blood glucose levels below the ideal range or decreasing at an alarming rate, i.e. decreasing by more than 4 mmol/L since last hour.

It should be mentioned that the protocol recommends one unit of insulin bolus for every 10 grams of carbohydrate consumed by the patient. However, not all protocols take carbohydrate consumption into account as patients do not eat when some protocols are in use, e.g. ICU protocols generally do not consider food intakes as patients do not eat in intensive care. The boluses are discarded in the continuous optimization process.

2.1.2. Glucose regulation for intensive care patients

GRIP² is an algorithm-based protocol for calculating suitable intravenous insulin infusion rate adjustments, originally intended for patients in ICU. In our research, GRIP is utilized in a more general non-ICU setting in which any in-patient in need of intravenous insulin infusion must be treated.

In GRIP, the recommended change in insulin infusion rate is determined by the mean insulin infusion rate of the last four hours ($\overline{I_{-4h}}$), the current blood glucose level (G_0), the target blood glucose level (G_{target} , set to 6.5 in GRIP), and the change in blood glucose level in the last four hours ($\Delta_{4h}G$). The blood glucose levels are measured in mmol/L. If there is less than four hour's worth of data, e.g. the infusion has just started, the infusion rate can be assumed to be zero before the starting point and the blood glucose level four hours ago can be assumed to be equal to the earliest measurement, and the output can then be calculated based on these assumptions. GRIP's output is the recommended insulin infusion rate change (ΔI) measured in U/h. This is primarily a recommendation for medical staff which may be overruled at any time according to medical judgement. The GRIP formula is shown in Eq. (1).

$$\Delta I = (1 + 0.25 \times \overline{I_{-4h}}) \times (0.2(G_0 - G_{\text{target}}) + 0.3 \times \Delta_{4h}G). \quad (1)$$

GRIP is essentially a linear combination of four input variables, i.e. $\overline{I_{-4h}}$, G_0 , G_{target} and $\Delta_{4h}G$, with G_{target} usually set to a constant in each session of intravenous insulin infusion therapy. In practice, further constraints are usually added to GRIP's recommendation, for example rounding large positive values of ΔI down to 1.5 U/h and rounding large recommended infusion rates down to 10 U/h.² In our experiments these constraints are treated as optional "post processing" of the recommendations that are specific to the hospital itself and for optimization purposes the focus is solely on the formula given in Eq. (1).

2.1.3. Other protocols

Glucommander is another simple and effective algorithm-based protocol.⁶ It can be presented as a set of linear models of the blood glucose level, each having a different slope, called a multiplier in the context of the protocol. The current multiplier is calculated using a function of the previous multiplier, the previous blood glucose level and the current blood glucose level. The recommended insulin infusion rate is

computed as the product of the current multiplier and the current blood glucose level offset by a constant. Lack of real-world data collected from Glucomander makes it hard to be benchmarked against other protocols using known methods,³ as it would be hard to reverse-engineer the multiplier at each recommendation.

Glucosafe is one of the more sophisticated protocols available.⁷ Glucosafe maintains a virtual model of the patient, i.e. modeling the state of the blood, the organs, etc., in order to calculate the optimal insulin infusion rate based on the virtual patient model. The sophistication of the protocol makes it not as good a candidate for AI optimization techniques used in our experiments.

According to the literature, some hospitals utilize an insulin infusion rate recommendation protocol, the core of which is a sigmoidal function of the blood glucose level.^{8,9} The sigmoidal protocol can be translated into a tabular protocol with textual instructions that is very similar to the protocol used at the Waikato DHB, with the blood glucose levels divided into ranges, the insulin infusion rate recommendations rounded to one decimal place, its various multipliers represented as a series of infusion scales, and the infusion rate scale in use adjusted by a set of textual instructions. Its simplicity, as well as its similarity to the Waikato DHB protocol, makes it an excellent protocol for comparison purposes during the validation process of the algorithmic protocols produced by continuous optimization.

2.2. Metaheuristics for continuous optimization

A continuous optimization problem is an optimization problem where the variables to be optimized are numeric, as opposed to discrete. The constants in an insulin infusion rate calculation protocol are all real numbers, which means continuous optimization algorithms are suitable for optimization of such protocols. The DEAP framework is used to facilitate our continuous optimization experiments.¹⁰ Two algorithms are used, i.e. Particle Swarm Optimization (namely, a specific type of PSO called Constriction Coefficient Particle Swarm Optimization)¹¹ and Covariance Matrix Adaption Evolution Strategy (CMA-ES).¹²

PSO explores the search space using a population of particles, and in each iteration the velocity of each particle is adjusted according to the known global and individual optimums in the search space. Bonyadi and Michalewicz¹¹ pointed out that the original PSO algorithm may result in infinite velocity of its particles, causing the particles to leave the reasonable search space. To prevent this behavior, they presented a variant of the PSO formula in their review, named Constriction Coefficient Particle Swarm Optimization (CCPSO), which is the version we used.

CMA-ES explores the search space by randomly generating a population of potential solutions per iteration based on an evolving distribution, composed of a mean, a covariance matrix and a set of standard deviations along the axes.¹² The mean and the covariance matrix are derived from the best individuals in the population. The step size of the optimization, i.e. the set of standard deviations, is controlled by how much the directions of the steps differ.

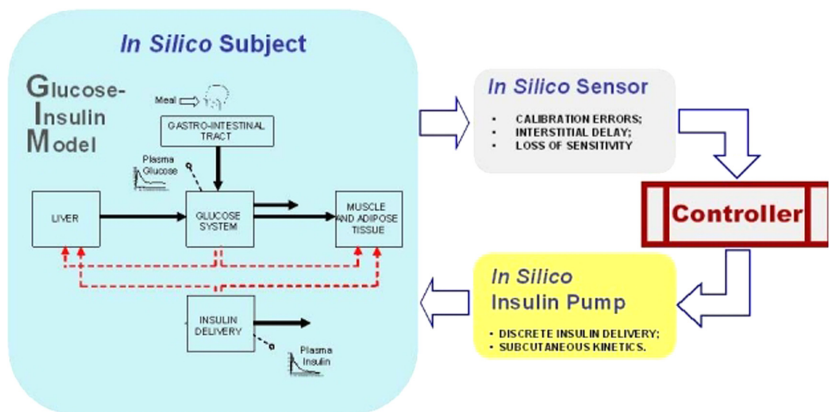
2.3. Simglucose

Simglucose is a python implementation of the FDA-approved UVa/Padova Simulator (2008 version)⁴ for research purpose only. It is implemented by Jinyu Xie. The online version is Simglucose v0.2.1 (2018). Source: Ref. 14, accessed on 14/01/2020.

Simglucose is able to simulate thirty virtual patients with Type 1 Diabetes, including 10 adults, 10 adolescents and 10 children. This is a valuable resource to our experiments as it can be used to verify the performance of the insulin infusion rate recommendation protocols produced by the continuous optimization experiments.

The virtual model of Simglucose is depicted in Fig. 3. Simglucose is designed to evaluate insulin administration protocols used by an Artificial Pancreas (AP), which measures blood glucose levels and administers insulin subcutaneously, as opposed to the intravenous treatment received by hospital in-patients. The AP is a device that is carried by a patient with Diabetes, constantly monitoring the patient's glucose level and administering insulin in order to keep the patient's glucose level stable and in range. The AP is very similar in concept to the Artificial Cardiac Pacemaker which is carried by patients with heart problems and regulates a patient's heart functions, while the AP regulates a patient's glucose–insulin balance. Simglucose simulates a virtual patient's blood glucose level, plasma insulin concentration, and the interaction between glucose and insulin, in addition to the relation between the subcutaneous system and the plasma balance.

In order to adapt the Simglucose system to simulate a hospital in-patient's intravenous insulin infusion treatment, where insulin is infused directly into the blood and the blood glucose level is measured from fingerstick blood, certain simulation components need to be replaced. As subcutaneous glucose levels are determined by blood glucose levels, the simulated blood glucose levels can be directly read as



Source: Ref. 13.

Fig. 3. The glucose–insulin model simulated by Simglucose for evaluating the performance an Artificial Pancreas insulin administration protocol.

measurements obtained from fingerstick blood. In the same vein, as insulin administered subcutaneously gets transferred into blood plasma, the transfer mechanism can be replaced with a protocol's recommended insulin infusion rate so that intravenous insulin therapy can be simulated. The modified system is crosschecked with a separate publication on intravenous insulin pharmacokinetics¹⁵ and the modified system configurations are confirmed to be valid.

3. Data Collection

The dataset used in our research was collected from the Waikato District Health Board, New Zealand. The original data was in paper form, which was recorded by nurses during the intravenous insulin infusion therapy sessions of hospital in-patients.

Data of 16 intravenous insulin infusion therapy sessions were collected, and these therapy sessions were received by 13 in-patients at the local hospital, three of whom received two therapy sessions each during different hospital admissions. In our experiments, all the 16 therapy sessions are treated as from different independent patients to avoid unnecessary complexity and produce a protocol that can make reasonable insulin infusion rate recommendations without any prior knowledge on the patient.

The data of each patient are in the form of a time series, where each time point stands for an insulin infusion rate adjustment and includes the following attributes:

- the date (DD/MM/YY),
- the time (hh:mm),
- the blood glucose levels measurement (mmol/L),
- the change in the blood glucose levels measurement since the last measurement (mmol/L),
- the insulin infusion scale used by the nurse when referring to the insulin infusion table,
- the insulin infusion rate decision (units/h) and
- the IV fluids and rate.

Some of the patients received insulin boluses in tandem with the intravenous insulin infusion. The type, dosage and time of each insulin bolus was obtained alongside the time series. However, the boluses are not taken into account during the continuous optimization process to find a competitive recommendation protocol, and the reasons are as follows:

- (1) GRIP serves as the starting point of the optimization, and GRIP does not utilize any insulin bolus, with the only insulin intakes being intravenous insulin infusion. This is presumably due to GRIP being designed as an ICU protocol, while the dataset for our experiments is collected from hospital in-patients. In addition, the creators of GRIP believed that insulin boluses may intensify an effect called

- insulin saturation.² Insulin saturation refers to the phenomenon where excessive insulin in a patient's plasma cannot be utilized by it receptors immediately, and may lead to increased hypoglycemia risks later on.¹⁶
- (2) At the local hospital, the reason for an insulin bolus is usually to counteract a recent carbohydrate intake, as it raises the blood glucose level temporarily. However, food intake data was not collected from the hospital archive, and is thus unavailable at the time of the optimization experiments. An additional point is that the food intake data may not be beneficial to the optimization nevertheless, as the effect of each carbohydrate intake on the blood glucose level is hard to quantify and will only make the optimization process harder to facilitate and potentially less accurate. Also, the starting point of optimization, GRIP, does not consider either carbohydrate intakes or boluses. Therefore, there is hardly a justifiable way of introducing carbohydrate influence.

According to the reasons above, it was decided that a reasonable course of action would be to optimize GRIP without considering either carbohydrate intakes or boluses, and the resulting protocol should run in tandem with a heuristic for insulin boluses, e.g. $\text{bolus (units)} = \frac{\text{carbohydrate (grams)}}{10}$ which is the current practice at the local hospital, if carbohydrate intakes are present.

The IV fluid was disregarded as well, because its effect, more specifically the effect of the dextrose in it on the blood glucose level, is hard to quantify. Another practical reason is that this data is not present on every sheet of written data and the amount of missing data is significant. It was deemed best to assume that the rules for IV fluid are similar enough in the scenarios of both GRIP and the local hospital.

Statistical analysis is performed on the data. There are 462 data points in the entire dataset, with each data point representing an insulin infusion rate adjustment.




Fig. 4. Blood glucose status of each patient.

The median number of adjustments per patient is 23, with the maximum number of adjustments being 88 and the minimum number being 7.

Figure 4 shows the blood glucose status of each patient during their insulin infusion session, i.e. the percentage of time with the patient's blood glucose level in the normal range, as well as above and below the normal range. On average, the patients' blood glucose levels are in the normal range for approximately 50% of the time, and it is believed that there is room for improvement, which our research aims to achieve.

4. Method

This section covers the methodology of our research, including the favorability evaluation method, different optimization strategies, our approach to optimization and the verification process.

4.1. Evaluation framework

Wong *et al.*³ proposed a methodology for evaluating algorithm-based insulin infusion rate recommendation protocols against certain historical practice, which is usually carried out using an algorithm-based protocol as well. The fundamental idea is to estimate the insulin infusion rate that *would have been* recommended by the protocol to be evaluated (e.g. GRIP), and how each of these hypothetical recommendations compares to each of the corresponding actual recommendations made by the protocol used in the historical practice (e.g. the method currently in use at the local hospital).

Historical data from a set of different patients over the course of their insulin infusion treatment are needed in order to facilitate the comparison. Ideally, each patient should have at least 12 hours' worth of data. Each patient's dataset must consist of a pair of time series, with one measuring the patient's blood glucose level at each infusion rate decision (i.e. at each time point), and the other recording the continuous insulin infusion rate at each of these decision points. Figure 5 depicts the insulin infusion treatment of one such patient. The time points are separated by approximately one hour in our data, which is consistent with the common intravenous insulin therapy practice of approximately hourly adjustments.

Based on the historical data, the algorithmic protocol to be evaluated is invoked, for each patient and each point in time, to calculate the insulin infusion rate that it *would have recommended*, had it been used at that point, given the past blood glucose levels and the amount of insulin already administered at the previous time points of the same insulin therapy session. Due to different algorithms being used in the historical practice and the protocol to be evaluated, the recommendations they make are most likely to differ. This is the key point of the evaluation.

In order to assess whether the recommendation made by the protocol to be evaluated is potentially better than that of the historical practice at a time t , the blood glucose level at the next decision point in the time series (i.e. at $t + 60$ minutes) needs to be examined. Essentially, if the future blood glucose level is above the normal range, then the higher insulin infusion rate recommendation is to be

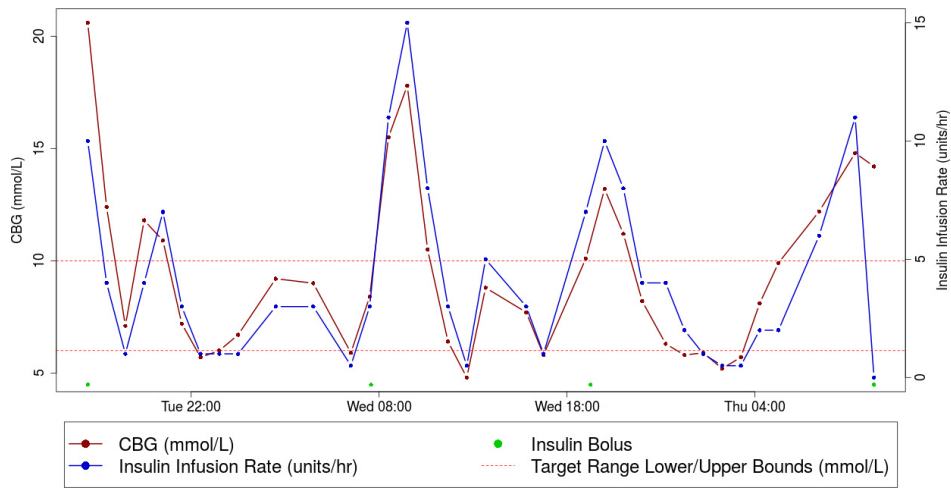


Fig. 5. (Color online) Example visualization of an insulin infusion session for a patient. Red horizontal lines indicate bounds for the normal range.

preferred, and consequently, as is the algorithmic protocol that would have made that recommendation. In other words, an above-range blood glucose level at the next time point indicates that the amount of insulin administered in the historical practice was insufficient to minimize hyperglycemia risk, and an algorithm that would have recommended a higher infusion rate could have performed better in keeping the blood glucose level in range. Conversely, if the future blood glucose level is below the normal range, then an algorithm recommending a lower insulin infusion rate is to be preferred as this would have reduced the hypoglycemia risk. If the recommendations made by both protocols are approximately the same, which in our experiments is defined as the difference being below a threshold of 10% of the historical recommendation, then they can be considered equivalent.

By aggregating the outcomes of all historical data points of all patients evaluated according to the principle above, one can estimate the overall performance of the algorithm to be evaluated compared to that of the algorithm used to produce the historical data. This aggregation of individual outcomes is referred to as the “favorability” of the algorithm.

Table 2 defines key symbols used in the favorability measurement. Algorithm 1 formalizes the favorability measurement for clarity. The algorithm depicts the steps required to calculate the favorability of a protocol relative to the historical practice at a single point in time t , with the protocol denoted by a vector β . This algorithm is essentially the same approach as described by Wong *et al.*³ but presented as pseudo code. In order to calculate the overall relative favorability of a protocol based on all available historical data, the algorithm needs to be applied to every consecutive point in time for every patient in the historical data and the outputs need to be summed.

Table 2. Definitions of terms and symbols used in the favorability measurement. Note that while time index t is shown explicitly, most of these terms can also be indexed by patient.

Term	Definition
$\beta = [\beta_0, \dots, \beta_{n-1}]$	numeric vector defining the strategy being optimized ($n = 4$ or $n = 8$ depending on the strategy)
t	time
bg_t	historical blood glucose level at time t
I_t	historical insulin infusion rate at time t
ΔI_t^β	infusion rate change suggested by the strategy being tested at time t
I_t^β	recommended insulin infusion rate change at time t
δ	threshold for determining whether two insulin rate adjustments are equal, set to 0.1 (10%) in this paper
r_{hypo}	threshold for hypoglycemia, set to 6.0 mmol/L
r_{hyper}	threshold for hyperglycemia, set to 8.0 or 10.0 mmol/L
$\text{fav}_t(\beta)$	favorability at time t , defined by Algorithm 1
λ	penalty size term for limiting magnitude of ΔI_t^β
$\text{obj}(\beta)$	the objective function to maximize during optimization

Algorithm 1. Function definition for computing relative favorability of a strategy β at time t for a specific patient. Constants are defined in Table 2.

Function $\text{fav}_t(\beta)$: $f \leftarrow \text{case_low}(I_t, I_t^\beta, bg_{t+1}) +$ $\text{case_high}(I_t, I_t^\beta, bg_{t+1}) +$ $\text{case_target}(I_t, I_t^\beta, bg_{t+1})$ return f;	Function $\text{case_high}(I_t, I_t^\beta, bg_{t+1})$: if $bg_{t+1} > r_{\text{hyper}}$ then if $I_t^\beta \geq I_t \cdot (1 + \delta)$ then return 1 end end return 0
Function $\text{case_low}(I_t, I_t^\beta, bg_{t+1})$: if $bg_{t+1} < r_{\text{hypo}}$ then if $I_t^\beta \leq I_t \cdot (1 - \delta)$ then return 1 end end return 0	Function $\text{case_target}(I_t, I_t^\beta, bg_{t+1})$: if $bg_{t+1} \geq r_{\text{hypo}}$ & $bg_{t+1} \leq r_{\text{hyper}}$ then if $I_t^\beta \leq I_t \cdot (1 - \delta)$ then return 1 end end return 0

Algorithm 1 is consistent with the framework published by Wong *et al.*³ However, it should be pointed out that the algorithm may have potential caveats, which are listed in what follows.

First of all, the algorithm treats the cases where the next blood glucose level is in range or below range essentially in the same way (i.e. when the next blood glucose measurement is on target or below target, the protocol recommending the lower

insulin infusion rate is considered more favorable). This is a potential weakness as the in-range and below-range cases are not differentiated, implying that only the upper bound of the target range (r_{hyper}) is significant in the favorability calculation, and the lower bound of the target range (r_{hypo}) serves no practical purpose. It can be argued, however, that a justifiable fix is hard to find. In clinical scenarios where hospital in-patients receive intravenous insulin infusion treatment, the risk of hypoglycemia from insulin over-dosing is much more significant and threatening to the patient than the risk of hyperglycemia from insulin under-dosing. Therefore, when the patient's blood glucose level is in range, clinicians should almost always prioritize avoiding insulin over-dosing rather than insulin under-dosing.

Second, the favorability of the algorithm to be evaluated is most likely an over-estimate. When the blood glucose level is out of range, the favorability measurement in itself cannot differentiate between a reasonable infusion rate change that returns the blood glucose levels back to normal (indeed favorable) and an extreme change that reduces blood glucose levels too much (in fact not favorable), and it will label both as favorable. For the favorability evaluation, there is an implicit assumption that all protocols being compared will recommend reasonable insulin doses at all time — however, this is *not* necessarily the case in our situation where we are starting our optimization from random protocols. To address and correct this issue, a novel regularization or penalty term is added to the optimization process, which leverages the connection between overestimations and extreme insulin rate changes. This is a contribution of our research to the favorability evaluation framework, as it introduces a way to gauge the overestimation in the favorability metric, and how to penalize it, which will be further elaborated in the following sections.

4.2. Strategies

The novel methodology utilized in our research is to generalize the GRIP formula in order to optimize it. The objective to be optimized is the relative favorability of the current protocol (i.e. generalized GRIP) compared with the historical practice (i.e. Waikato DHB protocol) on historical data (i.e. the processed dataset of 16 hospital in-patients), as described in Sec. 4.1.

For experimental comparison purposes, two baseline strategies are included, i.e. A and B. Strategy A, shown in Eq. (2), serves conceptually as the “No Change” baseline, and proposes no change to the current insulin infusion rate. In other words, the insulin infusion rate recommendation at a time point is the same as the historical insulin infusion rate at the previous time point, i.e. the previous insulin infusion rate decision. Although Strategy A is expected to perform poorly and be impractical, it is included as a “sanity check” for the other strategies.

$$\Delta I = 0. \quad (2)$$

The second baseline method, Strategy B, is the standard GRIP formula shown in Eq. (1). GRIP takes a single blood glucose target value instead of a blood glucose

range, i.e. G_{target} in Eq. (1). Originally, GRIP has a target value of 6.5 mmol/L which can be reasonably translated into a target range of 4–7.5 mmol/L. In addition to the original 6.5 mmol/L target value, a target value of 8 mmol/L is evaluated in Strategy B as well, as this target value is intended to represent the 6–10 mmol/L range, which is the target range of the local DHB practice and the favorability measurement of our research.

Next, two generalizations of the GRIP formula with optimizable parameters are defined. First of these, Strategy C replaces the constants in the original GRIP formula with variables to be optimized. The strategy is defined by Eq. (3). Strategy C contains four tunable parameters, each corresponding to a constant in the original GRIP protocol. To obtain the original GRIP formula, the four variables simply need to be assigned as follows: $\beta_0 = 1, \beta_1 = 0.25, \beta_2 = 0.2, \beta_3 = 0.3$.

$$\Delta I^\beta = (\beta_0 + \beta_1 \times \overline{I_{-4h}}) \times (\beta_2(G_0 - G_{\text{target}}) + \beta_3 \times \Delta_{4h}G). \quad (3)$$

Since Strategy C is relatively low-dimensional with only four β values to optimize, Strategy D is proposed, which is higher-dimensional, with greater flexibility in modeling insulin infusion rate recommendation protocols, and represents a larger space of potential protocols. Strategy D is defined by Eq. (4).

$$\begin{aligned} \Delta I^\beta = & \beta_0 + \beta_1 \times \overline{I_{-4h}} + \beta_2 \times (G_0 - G_{\text{target}}) + \beta_3 \times \Delta_{4h}G \\ & + \beta_4 \times \overline{I_{-4h}} \times (G_0 - G_{\text{target}}) + \beta_5 \times \overline{I_{-4h}} \times \Delta_{4h}G \\ & + \beta_6 \times (G_0 - G_{\text{target}}) \times \Delta_{4h}G \\ & + \beta_7 \times \Delta_{4h}G \times (G_0 - G_{\text{target}}) \times \overline{I_{-4h}}. \end{aligned} \quad (4)$$

Essentially, Strategy D assigns an optimizable parameter as a coefficient to each of the input arguments of the original GRIP formula (i.e. $\overline{I_{-4h}}$, $G_0 - G_{\text{target}}$ and $\Delta_{4h}G$) and all their interactions (e.g. $\Delta_{4h}G \times \overline{I_{-4h}}$). A tunable bias is also included as β_0 . Therefore, Strategy D has a total of eight tunable parameters. To obtain the original GRIP formula, the eight variables simply need to be assigned as follows: $\beta_0 = 0, \beta_1 = 0, \beta_2 = 0.2, \beta_3 = 0.3, \beta_4 = 0, \beta_5 = 0.075, \beta_6 = 0.05, \beta_7 = 0$.

It should be noted that G_0 and G_{target} are grouped together to represent the error between the current blood glucose levels and the target value. In practice, the target value is selected by a clinician according to the patient's physical condition, and is fixed during each insulin infusion therapy session, i.e. G_{target} is a constant during each session. Therefore, it makes sense to keep G_0 and G_{target} as one term when expanding the original GRIP formula into Strategy D, as doing so reduces the number of optimizable parameters without sacrificing any modeling capacity of the strategy. Additionally, the optimization process can always adjust the β values to fit the training data regardless of the value of G_{target} , and the significance of G_{target} is in providing clinicians with an adjustable target value/range for each insulin therapy session.

4.3. Approach to optimising the strategies

Based on the description of the metaheuristic search algorithms in Sec. 2.2, this section provides further details on the individuals in the population and the fitness function in the context of our research.

Although PSO and CMA-ES are different algorithms, the list of optimizable parameters can be implemented in the same straightforward fashion in both as a vector that represents the space of potential solutions. For Strategy C, an individual is a vector containing β_0 to β_3 , which conforms with the four-dimensional search space of Strategy C, and each individual contains its own β_0 to β_3 values that get improved during the optimization process. It is the same with Strategy D, but with β_0 to β_7 representing the eight-dimensional search space of the strategy.

The objective function, whose return value is to be maximized for optimization, calculates the overall fitness of an individual. An individual's recommended insulin infusion rate change at each time point can be calculated from the the strategy's formula with the β vector of the individual, and the $\overline{I_{-4h}}$, $G_0 - G_{\text{target}}$ and $\Delta_{4h}G$ values in the historical data. Then the individual's recommended insulin infusion rate at each time point can be calculated by adding its recommended change and the historical insulin infusion rate at the previous time point. Finally, the individual's favorability at each time point is determined by comparing its recommended insulin infusion rate against the historical rate in the context of the blood glucose level at the next time point (i.e. high glucose levels favor higher insulin infusion rates and low glucose levels favor lower insulin infusion rates, refer also to Algorithm 1). The overall favorability score of a potential protocol is the sum of the protocol's fitness values at all the time points of all the patients' insulin infusion therapy sessions. As explained in Sec. 4.1, a higher overall favorability score indicates better performance of a protocol. Therefore, better configurations for Strategy C or D can be found by a continuous optimization algorithm by trying to maximize its population's fitness scores. It should be pointed out that for each session the first and last time points should be discarded for favorability evaluation because the first time point lacks a previous historical insulin infusion rate and the last time point lacks a next blood glucose measurement.

The favorability score alone is inadequate in representing a reasonable fitness function, however. The favorability score has a tendency to increase the absolute value of a protocol's recommended insulin infusion rate change. The reason is speculated to be that, under the sole optimization pressure from the favorability measurement, the continuous optimization process can only improve recommendations that are unfavorable, with no idea of how to improve favorable recommendations, some of which may be excessive and not truly beneficial. This phenomenon is analogous to over-fitting in Neural Networks, where the system satisfies the evaluation metric (e.g. the favorability metric) at the expense of practicality (e.g. being a reasonable and useful protocol). As discussed in Sec. 4.1, the strong positive correlation between the absolute value of a recommended change and the risk of

overestimation calls for a penalty term to mitigate this effect. It is deemed appropriate to use, as a penalty term, the sum of the squared values of the insulin infusion rate recommendations, i.e. at each time point a protocol is not only rewarded according to its favorability, but also penalized according to the magnitude of its recommended change. The penalty term is scaled by a user-specified parameter, as a smaller multiplier emphasizes the favorability metric while a larger multiplier focuses more on the penalty term, with the former encouraging more aggressive favorable recommendations and the latter encouraging more moderate recommendations. The good “middle ground” of this trade-off needs to be determined by the user by specifying an effective penalty scale, in order to produce a protocol which makes favorable recommendations with as little excess as possible.

The return value of the fitness function is the overall favorability score, minus the penalty term, as shown in Eq. (5).

$$\text{obj}(\beta) = \sum_{p \in \text{patients}} \left(\sum_t \text{fav}_{t,p}(\beta) - \lambda \sum_t ((\Delta I_{t,p}^\beta)^2) \right). \quad (5)$$

The equation shows that, given a protocol parameterized with a vector of tunable variables β , the function sums over all the time points of all the patients. The weighting of the penalty term is determined by a user-specified parameter λ .

4.4. Post-optimization verification

The resulting protocols obtained from the continuous optimization process can be validated with Simglucose. Simglucose provides a good framework for evaluating different insulin administration protocols, and with some modifications can perform reasonably good simulations of intravenous insulin infusion therapy, as described in Sec. 2.3. Simglucose performs its simulations using virtual patients, whose data have never been exposed to the continuous optimization process. This is a key point to the reliability of the validation outcome as all the protocols evaluated are equally blind to the blood glucose behaviors of the simulated virtual patients.

Ten virtual adult patients with Type 1 Diabetes in the Simglucose implementation are used in the simulated validation and each simulation represents 240 h of intravenous insulin therapy (i.e. 10 days of continuous insulin infusion). Among the 30 virtual patients, 10 adults are closest to the 16 real patients in the dataset in terms of physical condition, while 10 adolescents deviate from the dataset in physical condition such as body weight and insulin sensitivity and 10 children deviate even further. The 10 virtual adults are the most similar to the intended recipients of the protocols, and therefore used in our validation process. The total number of runs is 100 as 10 runs are performed on each patient.

Simglucose simulates carbohydrate intakes (i.e. meals), ultimately these are included in the simulations, which means the virtual patients consume food as a Type 1 Diabetes patient would during their 10-day insulin therapy sessions. The decision was made because simulations without meals turn out to be unchallenging and any

reasonable protocol can stabilize the patients' blood glucose levels with utter ease. Another reason is that meals introduce more randomness into the simulations as blood glucose reading errors contain very limited randomness.

Three straightforward insulin bolus strategies are used in the simulations to assist the infusion protocols in order to measure the performance of the protocols as well as the bolus strategies. The three bolus strategies are as follows:

- No bolus whatsoever, relying entirely on the insulin infusion process for blood glucose regulation after eating.
- Each time the patient consumes n grams of carbohydrate, the patient will be given $\frac{n}{10}$ units of insulin bolus directly into the bloodstream in 15 min, which is the current local DHB practice in insulin therapy for hospital in-patients.
- Each time the patient consumes n grams of carbohydrate, the patient will be given $\frac{n}{5}$ units of insulin bolus directly into the bloodstream in 15 min, which is a more proactive approach in avoiding after-meal hyperglycemia. This strategy is suggested for investigation by the local DHB Diabetes expert, as mild after-meal hyperglycemia is often observed among the hospital in-patients undergoing insulin therapy using the current bolus strategy.

Rules are obtained from a Diabetes expert at the Waikato DHB to distinguish an acceptable simulation run from an unacceptable one. All these rules must be satisfied for a simulation run to be acceptable:

- (1) Normoglycemia (4–10 mmol/L inclusive) must be maintained for at least 50% of the time in an insulin therapy session
- (2) Significant hyperglycemia (over 14 mmol/L) must not occur for more than 10% of the time in an insulin therapy session
- (3) Hypoglycemia (Below 4 mmol/L) must not occur for more than 1% of the time in an insulin therapy session

5. Results

The experimental results are covered here.

5.1. Optimization experiment results

This subsection analyzes the results of the optimization experiments.

5.1.1. Result overview

The entire 16-patient dataset contains 430 data points suitable for the favorability measurement, excluding the first and last time points of each patient as described in Sec. 4.3. Therefore, the theoretical maximum favorability that can be achieved by a protocol is 430, which means that such a protocol would be more favorable than the Waikato DHB protocol at every available time point in the dataset.

Strategy A is evaluated to have an overall favorability of 90, which is very poor and fully anticipated, as Strategy A is included as a baseline algorithm analogous to the no-change algorithm.

Strategy B is evaluated with two different target values, 6.5 mmol/L for a target range of 4–7.5 mmol/L and 8 mmol/L for a target range of 6–10 mmol/L. Strategy B with 6.5 mmol/L has a favorability of 202, while Strategy B with 8 mmol/L has a favorability of 215. Note that here Strategy B refers to the GRIP formula (Eq. (1)) but does not include any further constraints involved in real life scenarios where GRIP is utilized. Indeed, these further constraints should be considered post-processing, a great portion of which varies from hospital to hospital, as well as from patient to patient. Given the task of optimising the GRIP formula (Strategy C and Strategy D), it is best to evaluate the GRIP formula solely as a benchmark protocol.

The continuous optimization experiment results of Strategy C and Strategy D are shown in Tables 3–6, also visualized in Figs. 6–9. During each experiment, the population is optimized for 100 iterations. For either Strategy C or D, 36 configurations are explored, involving the PSO and CMA-ES algorithms, population sizes of 80, 160 and 320, as well as penalty scale λ values of 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} and 10^{-1} (i.e. $2(\text{algorithms}) \cdot 3(\text{population_sizes}) \cdot 6(\text{penalty_scales}) = 36(\text{configurations})$). Each configuration is evaluated through 10 runs, each of which produces a potential protocol at the end of the continuous optimization process, and Tables 3–6 show the statistics of the 10 runs, i.e. the best and median values of the favorability measurement and the penalty term before scaling.

The tables indicate that Strategy D clearly outperforms Strategy C, as Strategy D achieves significantly greater favorability, compared with Strategy C given the same configuration, without the penalty growing significantly. The greatest favorability achieved by Strategy D is 309 while Strategy C achieved a maximum of 280.

Table 3. Favorability results for strategy C.

Best fav (Median fav)		λ					
		10^{-6}	10^{-5}	10^{-4}	10^{-3}	10^{-2}	10^{-1}
PSO	80	279 (278)	277 (276)	271 (271)	262 (259)	229 (223)	184 (175.5)
		279 (278)	277 (275.5)	272 (271)	264 (262.5)	231 (223.5)	185 (177)
	160	279 (279)	277 (276.5)	274 (271)	264 (261.5)	233 (226.5)	184 (179)
		279 (279)	277 (276.5)	274 (271)	264 (261.5)	233 (226.5)	184 (179)
	320	279 (279)	277 (276.5)	274 (271)	264 (261.5)	233 (226.5)	184 (179)
		279 (279)	277 (276.5)	274 (271)	264 (261.5)	233 (226.5)	184 (179)
CMA-ES	80	280 (278.5)	277 (276.5)	271 (271)	260 (260)	228 (227.5)	180 (180)
		279 (279)	277 (277)	271 (271)	260 (260)	233 (228)	180 (180)
	160	279 (279)	277 (277)	271 (271)	260 (260)	233 (228)	180 (180)
		279 (279)	277 (277)	271 (271)	260 (260)	233 (228)	180 (180)
	320	280 (279)	277 (277)	271 (271)	260 (260)	231 (226)	180 (180)
		280 (279)	277 (277)	271 (271)	260 (260)	231 (226)	180 (180)

Table 4. Penalty sum results for strategy C.

Min sum (Median sum)		λ					
		10^{-6}	10^{-5}	10^{-4}	10^{-3}	10^{-2}	10^{-1}
PSO	80	7.4e5 (8.3e5)	1.2e5 (2.6e5)	3.3e4 (6.4e4)	9.3e3 (1.2e4)	1.7e3 (2.1e3)	1.8e2 (2.6e2)
	160	3.0e5 (7.9e5)	1.1e5 (1.7e5)	3.2e4 (6.4e4)	1.1e4 (1.5e4)	1.7e3 (2.0e3)	2.1e2 (2.7e2)
	320	3.6e5 (7.8e5)	1.2e5 (2.6e5)	4.3e4 (6.0e4)	9.6e3 (1.4e4)	1.7e3 (2.3e3)	2.3e2 (2.8e2)
CMA-ES	80	7.0e5 (1.2e6)	1.2e5 (2.4e5)	5.6e4 (5.6e4)	1.1e4 (1.1e4)	2.0e3 (2.3e3)	2.5e2 (2.5e2)
	160	7.0e5 (7.3e5)	1.9e5 (2.5e5)	5.6e4 (5.6e4)	1.1e4 (1.1e4)	1.9e3 (2.3e3)	2.5e2 (2.5e2)
	320	7.0e5 (7.6e5)	1.1e5 (2.5e5)	5.6e4 (5.6e4)	1.1e4 (1.1e4)	1.9e3 (2.1e3)	2.5e2 (2.5e2)

Table 5. Favorability results for strategy D.

Best fav (Median fav)		λ					
		10^{-6}	10^{-5}	10^{-4}	10^{-3}	10^{-2}	10^{-1}
PSO	80	303 (300)	298 (294.5)	290 (282.5)	275 (258)	253 (209)	219 (156.5)
	160	302 (299.5)	300 (297)	300 (286.5)	286 (268)	264 (203)	192 (156.5)
	320	305 (302)	301 (297)	293 (290.5)	277 (264)	263 (229)	191 (162.5)
CMA-ES	80	309 (309)	305 (305)	304 (304)	290 (287.5)	248 (243)	202 (194)
	160	309 (309)	306 (305)	305 (304)	290 (290)	248 (248)	200 (194)
	320	309 (309)	306 (305)	304 (301)	290 (290)	248 (248)	202 (200.5)

Focusing on Table 5, i.e. the favorability results of Strategy D, it can be seen that CMA-ES consistently outperforms PSO, except when the penalty scale λ is relatively large and arguably the favorability score becomes less than desirable. On the grounds that PSO may require more iterations to achieve the same level of performance as CMA-ES, the Strategy D experiments with the PSO configurations are re-run with 1,000 iterations instead of the original 100. However, no significant improvement is observed of the 1,000-iteration results. Therefore, it can be concluded that CMA-ES is the superior metaheuristic search algorithm for this specific problem.

By examining Tables 4 and 6, it can be seen that the negative correlation between the penalty scale λ and the unscaled penalty sum is clear. For example, focusing on the CMA-ES results of Strategy D, a small λ value such as 10^{-6} leads to very large

Table 6. Penalty sum results for strategy D.

Min sum (Median sum)		λ					
		10^{-6}	10^{-5}	10^{-4}	10^{-3}	10^{-2}	10^{-1}
PSO	80	3.5e5 (9.3e5)	1.9e5 (3.1e5)	4.5e4 (9.9e4)	6.3e3 (3.2e4)	4.6e3 (8.9e3)	1.2e3 (4.3e3)
	160	3.4e5 (6.4e5)	2.0e5 (3.8e5)	4.2e4 (8.7e4)	1.7e4 (2.4e4)	1.7e3 (6.6e3)	1.0e3 (4.0e3)
	320	3.4e5 (1.0e6)	1.7e5 (2.5e5)	3.2e4 (1.0e5)	1.2e4 (2.2e4)	3.3e3 (7.5e3)	8.8e2 (4.4e3)
CMA-ES	80	1.3e6 (1.5e6)	2.3e5 (2.7e5)	7.0e4 (7.7e4)	9.3e3 (1.2e4)	1.7e3 (2.1e3)	1.9e2 (2.1e2)
	160	1.5e6 (1.5e6)	1.2e5 (2.4e5)	6.5e4 (7.6e4)	1.2e4 (1.2e4)	2.0e3 (2.4e3)	1.8e2 (2.1e2)
	320	1.5e6 (1.5e6)	1.2e5 (2.3e5)	6.5e4 (7.1e4)	1.2e4 (1.2e4)	2.4e3 (2.4e3)	1.8e2 (2.2e2)

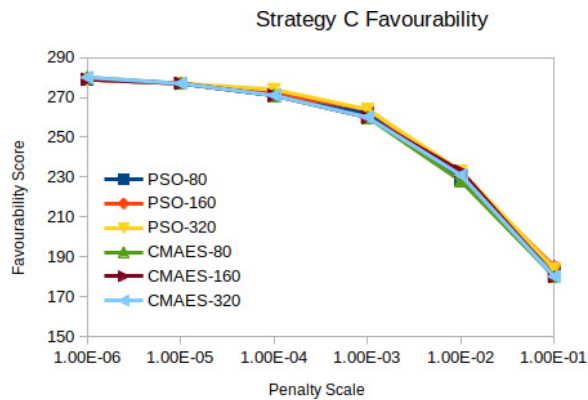


Fig. 6. Strategy C favorability results.

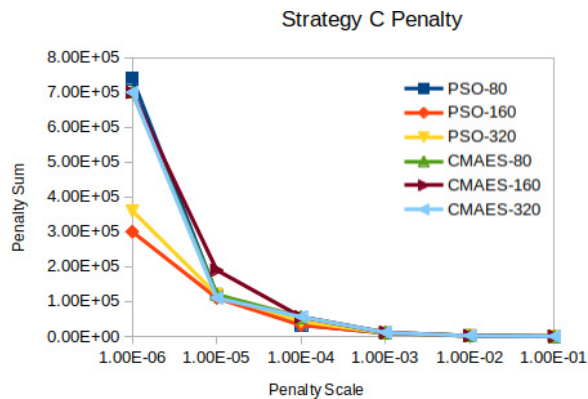


Fig. 7. Strategy C penalty results.

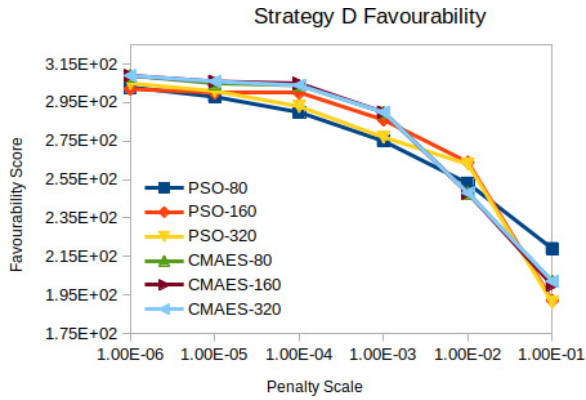


Fig. 8. Strategy D favorability results.




Fig. 9. Strategy D penalty results.

unscaled penalty sums like 1.5×10^6 , which indicates that the recommended insulin infusion rate changes ΔI are likely to be very large in its absolute value, giving the protocol a high favorability score on paper but making it potentially impractical. On the other hand, a large λ value such as 0.1 leads to small unscaled penalty sums like 1.8×10^2 , which suggests that the recommended insulin infusion rate changes ΔI are likely to be small, however, the favorability score may reach merely 200 in these configurations, making the optimization result even slightly less favorable than the original GRIP formula (which achieved a favorability score of 202 or 215 depending on the selected target value).

It can therefore be pointed out that the selection of a single optimal protocol requires a reasonable trade-off between the total favorability score and the penalty sum. As discussed above, extreme λ values, large or small, do not produce optimal protocols. Examining Strategy D and using Tables 5 and 6, it can be seen that the favorability decreases relatively slowly for λ values from 10^{-6} to 10^{-3} while the

penalty sum decreases significantly in the same range. However, a drastic drop in favorability begins at $\lambda = 0.01$ and making results in this range far less desirable than the previous range. It can therefore be argued that the λ value of 10^{-3} is the optimal amongst the six λ values.

It can be noted in addition that CMA-ES is relatively robust in terms of its population size, as no clear performance gain or loss in terms of favorability or penalty is observed relative to the population size in the CMA-ES experiments. It can therefore be speculated that all nontrivial population sizes produce similar results with CMA-ES, as the performance saturation in terms of population size can be attributed to the optimum in the search space already being sufficiently explored, therefore increasing the population size further does not grant improved performance.

5.1.2. Protocol analysis

One resulting protocol, which is deemed optimal of the experimental output, is given for further examination. The protocol comes from the following configuration:

- Strategy D,
- CMA-ES as the metaheuristic search algorithm,
- Population size of 320 and
- Penalty scale λ value of 10^{-3} .

Referring to Eq. (4), which is the formula for Strategy D, the exact values of the tunable parameters are shown in Table 7. The assignment of these variables results in a novel protocol shown in Eq. (6).

$$\begin{aligned} \Delta I^\beta = & -4.140 + 0.354 \times \overline{I_{-4h}} + 1.600 \times (G_0 - G_{\text{target}}) + 0.141 \\ & \times \Delta_{4h}G - 0.059 \times \overline{I_{-4h}} \times (G_0 - G_{\text{target}}) - 0.034 \times \overline{I_{-4h}} \\ & \times \Delta_{4h}G - 0.027 \times (G_0 - G_{\text{target}}) \times \Delta_{4h}G + 0.009 \\ & \times \Delta_{4h}G \times (G_0 - G_{\text{target}}) \times \overline{I_{-4h}}. \end{aligned} \quad (6)$$

Table 8 shows the protocol's recommendations for nine representative examples from the data. From this table, we can see that most of the recommendations seem

Table 7. Optimized variables of strategy D.

β_0	-4.140
β_1	0.354
β_2	1.600
β_3	0.141
β_4	-0.059
β_5	-0.034
β_6	-0.027
β_7	0.009

Table 8. Sample recommendations from the dataset.

Case	$\overline{I_{-4h}}$	G_0	$\Delta_{4h}G$	Rec. ΔI^β
1	3.7	11.8	-8.8	4.1
2	4.3	6.3	-2.1	-2.8
3	8.2	15.0	-6.2	6.7
4	5.3	13.7	0.2	7.0
5	3.0	9.9	2.8	1.9
6	1.6	12.4	3.0	5.3
7	8.5	14.7	0.9	8.1
8	2.0	11.8	1.4	4.5
9	3.5	4.2	-4.4	-6.2

reasonable. For example, in Case 1, where blood glucose levels are initially well above the 8 mmol/L target but decreasing over the past four-hour period, moderate increases in insulin rate are recommended. Case 2, on the other hand, represents a patient who is already below target but has been receiving some insulin over the past four hours. In this case, moderate negative adjustments are recommended. Case 4 shows a patient with high initial glucose levels that have not changed much in the past four hours, and therefore significant increments in insulin pump rate are recommended. Finally, Case 9 shows a patient at significant risk of severe hypoglycemia following a significant drops in glucose level. In this case, a large negative adjustment is recommended which effectively sets the insulin pump rate to zero (since negative pump rates do not make sense). The remaining cases show similar sensible recommendations being made.

Figure 10 is a visualization of the same strategy used to produce the recommendations in Table 8. To produce the visualization, we fixed values of $\overline{I_{-4h}}$ and $\Delta_{4h}G$ to different constants and plotted the resulting linear relationships between G_0 and ΔI . Note that the model itself is not linear in practice, since $\overline{I_{-4h}}$ and $\Delta_{4h}G$ would be constantly changing over time. The true model is a curved surface in four dimensions, nonetheless, this type of visualization enables us to see some of the model's properties.

The $\overline{I_{-4h}}$ and $\Delta_{4h}G$ inputs both appear to influence the slope of the $G_0/\Delta I$ relationship. In particular, negative values of $\Delta_{4h}G$ (indicating a decrease in blood glucose levels over the past four hours) and large values of $\overline{I_{-4h}}$ (indicating high insulin infusion rates over the past four hours) result in a gentler slope and therefore smaller rate increments when an increase is recommended. The figure shows that when $\overline{I_{-4h}}$ is set to 5.0, the protocol recommends no change when G_0 is approximately 8.0, regardless of $\Delta_{4h}G$. This indicates that the protocol considers the scenario where the blood glucose level is around 8.0 mmol/L and the average infusion rate has been around 5.0 U/h to be a relatively stable state, and recommends keeping the same insulin infusion rate until there are further glucose changes. For the plot, we




Fig. 10. Visualization of one optimized version of strategy D.

selected $\Delta_{4h}G$ values of 5.0 and -5.0 , since they are “typical” values for this variable in the data, while -20.0 and 15.0 are rare extreme values.

5.2. Verification experiment results

Table 9 summarizes the simulation results. It should be noted that both GRIP and Strategy D are given a target of 8 mmol/L, as well as GRIP’s practical constraints, i.e. $\Delta I \leq 1.5$ mmol/L and $I \leq 10$ mmol/L.

The Strategy D protocol achieved competitive results. It can be reasoned that to make a good formula-based protocol, both a good formula and reasonable engineering are necessary. However, an intriguing find is that when the formula was developed it was blind to the engineered features that would be added later (i.e. in the continuous optimization phase Strategy D was a pure formula without any manually defined constraints), yet it achieves top-level performance with common-sense constraints. This finding supports the viability of this metaheuristic approach, where interpretable models are developed through optimization and only modest

Table 9. Number of acceptable runs achieved by each infusion protocol.

	DHB	Sigmoid	GRIP	Strategy D
Standard bolus	46	29	2	54
Double bolus	31	21	10	51
No bolus	10	0	2	2

engineering is required to modify the models to make them suitable for problems in specific domains.

The table also shows that the protocols perform significantly worse without insulin boluses. Therefore, it is reasonable to say that in scenarios where the patient consumes carbohydrate, either insulin boluses are required or efforts need to go into making the continuous insulin infusion protocol handle carbohydrate intakes explicitly (i.e. increasing the insulin infusion rate according to the amount of food consumed), because trying to use the protocol to implicitly handle food intake (i.e. waiting for the protocol to recommend more insulin after the blood glucose level has already risen) tends to lead to more glucose fluctuations both above and below range, and ultimately leads to ineffective glucose control.

The protocols also tend to perform a bit better with the standard bolus scheme currently employed at the local DHB compared with the double-bolus scheme. It is speculated that a plain double-bolus scheme increases hypoglycemia risk in some patients who might have a higher insulin sensitivity. It remains to be seen if a potentially more complex bolus scheme, e.g. a scheme that recommends different amounts of bolus insulin for patients in different health conditions, can lead to significantly better glucose control.

6. Conclusion

Our research utilizes metaheuristic search algorithms to optimize GRIP with respect to historical data recorded by the local DHB, and the resulting protocols are validated with Simglucose and show superior performance when compared with GRIP and the local DHB protocol. In the end, the strategy D models are fairly optimized and extended to have become very different from the original GRIP formula, which is the starting point of its optimization.

Our research demonstrates the usefulness of Artificial Intelligence techniques that produce interpretable models. Such interpretable models can be understood and reasoned, and this is very important in domains where the behavior of software and tools need to be fully known, e.g. medical informatics. The interpretability also makes new ways of manipulating the model easier. Therefore, interpretable models combine nicely with elegant engineering, which makes it easy to do quick prototyping with such models.

The combination of the metaheuristic search algorithms and the favorability evaluation framework prove to work well and produce results of value, which is further verified by the simulations facilitated by Simglucose. The overall paradigm of optimizing on historical data and validating with simulations (or potentially the other way round, optimizing on simulations and validating with historical data) turns out to be efficient and robust for research endeavors. In this paradigm, only one real-world dataset is required, and concerns over the overlap between optimization and validation data should be minimum, as the two processes utilize data from very

different sources (with the real-world data being more statistical and the simulations being more theoretical).

Outstanding Strategy D protocols are found through simulated validation, and it is recommended that these protocols be further examined by medical experts in the diabetes field. If these examinations return positive results, further engineering on these protocols through collaborations between medical and computer science experts may produce optimal insulin rate protocols appropriate for clinical trials.

For future work, there are other optimization techniques that can be applied to our research problem, such as Genetic Programming,¹⁷ Neural Networks that produce interpretable models,¹⁸ etc. It is also an interesting idea to try to improve the strategy's response to meals, as GRIP does not take meals into account likely due to it being developed with ICU settings in mind. More explicit questions would be what formula to use for the protocol's response to meals and what kind of dataset is needed in order to optimize such a formula. The favorability framework's response to normoglycemia could also be improved, as currently it responds to normoglycemia in the same way as it does to hypoglycemia.

In addition, there are other protocols that may serve as good starting points like GRIP in this paradigm. Examples include PID-based controllers,¹⁹ mathematical-function-based strategies like the sigmoidal protocol,⁸ and instruction-and-table-based protocols such as the one used at the Waikato DHB.

Lastly, our research produces protocols that suggest insulin infusion rate changes but not the initial insulin infusion rate when a therapy session starts. Simple guidelines included in the current local DHB practice should suffice in providing an initial rate. It is nearly impossible for any protocol to come up with a perfect initial rate due to the great uncertainty involved at the start of a session. For a protocol, correctly adjusting the infusion rate as a session continues is far more important and feasible than merely coming up with a perfect rate initially, as long as its initial recommendations are reasonable. Nevertheless, it remains an interesting topic to discover strategies that could give better starts to insulin infusion therapy sessions.

References

1. K. Kaul, J. M. Tarr, S. I. Ahmad, E. M. Kohner and R. Chibber, *Introduction to Diabetes Mellitus* (Springer, New York, NY, 2013), pp. 1–11.
2. M. Vogelzang, F. Zijlstra and M. W. Nijsten, Design and implementation of GRIP: A computerized glucose control system at a surgical intensive care unit, *BMC Med. Inf. Decision Making* **5**(1) (2005) 38.
3. A. F. Wong, U. Pielmeier, P. J. Haug, S. Andreassen and A. H. Morris, An in silico method to identify computer-based protocols worthy of clinical study: An insulin infusion protocol use case, *J. Am. Med. Inf. Assoc.* **23**(2) (2016) 283–288.
4. B. P. Kovatchev, M. Breton, C. Dalla Man and C. Cobelli, In silico preclinical trials: A proof of concept in closed-loop control of type 1 diabetes, *J. Diabetes Sci. Technol.* **3**(1) (2009) 44–55, PMID: 19444330.

5. H. Wang, L. Chepulis, R. G. Paul and M. Mayo, Metaheuristics for discovering favorable continuous intravenous insulin rate protocols from historical patient data, *Intelligent Information and Database Systems* (Springer International Publishing, 2020), pp. 157–169.
6. P. C. Davidson, R. D. Steed and B. W. Bode, Glucommander: A computer-directed intravenous insulin system shown to be safe, simple, and effective in 120,618 h of operation, *Diabetes Care* **28**(10) (2005) 2418–2423.
7. U. Pielmeier, S. Andreassen, B. Juliussen, J. G. Chase, B. S. Nielsen and P. Haure, The glucosafe system for tight glycemic control in critical care: A pilot evaluation study, *J. Critical Care* **25**(1) (2010) 97–104.
8. R. Devi, T. Zohra, B. S. Howard and S. S. Braithwaite, Target attainment through algorithm design during intravenous insulin infusion, *Diabetes Technol. Therapeutics* **16**(4) (2014) 208–218.
9. S. Shapiro Braithwaite, L. P. Clark, T. Idrees, F. Qureshi and O. T. Soetan, Hypoglycemia prevention by algorithm design during intravenous insulin infusion, *Current Diabetes Rep.* **18**(5) (2018) 26.
10. F.-A. Fortin, F.-M. De Rainville, M.-A. Gardner, M. Parizeau and C. Gagné, DEAP: Evolutionary algorithms made easy, *J. Mach. Learn. Res.* **13** (2012) 2171–2175.
11. M. R. Bonyadi and Z. Michalewicz, Particle swarm optimization for single objective continuous space problems: A review, *Evolutionary Comput.* **25**(1) (2017) 1–54.
12. N. Hansen, The CMA evolution strategy: A tutorial, *CoRR*, abs/1604.00772.
13. Diabetes Technology Society, <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2681269/figure/fig1/>.
14. J. Xie, <https://github.com/jxx123/simglucose>.
15. B. Hipszer, J. Joseph and M. Kam, Pharmacokinetics of intravenous insulin delivery in humans with type 1 diabetes, *Diabet. Technol. Therapeutics* **7**(1) (2005) 83–93, PMID: 15738706.
16. S. Andreassen, U. Pielmeier and J. Chase, Receptor-based models of insulin saturation dynamics, in *Proceedings of the Sixth IASTED International Conference on Biomedical Engineering* (ACTA Press, Innsbruck, Austria, 2008), pp. 182–186.
17. A. M. Turing, Computing machinery and intelligence, *Mind* **49** (1950) 433–460.
18. S. S. Sahoo, C. H. Lampert and G. Martius, Learning equations for extrapolation and control, *CoRR* abs/1806.07259.
19. G. M. Steil, D. Deiss, J. Shih, B. Buckingham, S. Weinzimer and M. S. D. Agus, Intensive care unit insulin delivery algorithms: Why so many? How to choose? *J. Diabetes Sci. Technol.* **3**(1) (2009) 125–140.