UNIVERSITY OF PADOVA

DEPARTMENT OF INFORMATION ENGINEERING
Ph.D. Course in Information Engineering
Curriculum: Bioengineering
Series: XXXV

Machine Learning Based Techniques for the Design of Personalized Insulin Bolus Calculators in Type 1 Diabetes Therapy

Course director:
Andrea Neviani

Advisor:
Andrea Facchinetti

Co-advisors:

Giovanni Sparacino
Giacomo Cappon

Ph.D. Candidate:
Giulia Noaro

September 2022

Abstract

In recent years, the incidence and prevalence of type 1 diabetes (T1D) are increasing worldwide. In addition to the economic burden related to T1D, the management and treatment of such a disease require lots of effort from those people who are affected, as their body is no longer able to produce insulin, one of the key hormones in blood glucose (BG) regulation. The lack of endogenous insulin production results in elevated BG levels and, in particular, in hyperglycemia, a condition that can lead to several long-term cardiovascular complications, such as retinopathy and nephropathy. Therefore, people affected by T1D need lifelong therapy, which relies on exogenous insulin administrations. However, daily management of T1D significantly impacts on patient's quality of life, due to the number of tasks required to achieve proper glucose level regulation. As a matter of fact, one of the major obstacles to optimal glycemic control is represented by the estimation of a correct prandial insulin dose, which is injected to counteract the BG excursion following a meal. Indeed, an accurate mealtime insulin dosing in T1D therapy is crucial to avoid postprandial hypo- or hyperglycemic events, caused by an over- or under-dosage respectively. According to the recommended guidelines for T1D management, the mealtime insulin amount should be calculated following an empirical standard formula (SF), which could lead to a suboptimal dosage due to several reasons, including, above all, the inability of accounting for relevant information related to the glucose dynamics and a lack of individualization.

Prandial insulin estimation is a highly patient-dependent task which should be specifically tailored to the individuals' mealtime condition, not only by integrating relevant personalized parameters but also by adjusting the dose based on the current BG trend. Such information on BG dynamics and, in particular, its rate of change is provided, in real-time, by continuous glucose monitoring (CGM) sensors, minimally invasive devices that are becoming a key element in T1D therapy. The real-time availability of information on glucose dynamics provided by CGM systems, along with the possibility of leveraging smart insulin delivery devices, which could potentially integrate a novel dosing technique, fostered the development of new approaches to adjust the SF amount according to the glucose information provided by these sensors. However, the derivation of the proposed state-of-art approaches aimed at correcting the SF has mainly been empirical, suggesting that there would be room for improvement should a systematic modelling methodology be adopted.

Therefore, the work presented in this thesis aims at proposing effective and personalized mealtime insulin dosing techniques, which take into account both the CGM-derived information and the specific mealtime status of the individual, to optimize such a dosage, by leveraging machine learning and reinforcement learning algorithms. As a starting point, an assessment of the current state-of-art methods for insulin dose adjustment, which take into account the CGM rate-of-change information, was performed through an ISCT, together with the design of a novel, yet empirical, method based on the literature. This analysis pointed out that there is no literature method which outperforms the others in terms of postprandial glycemic control, highlighting the need for insulin dosing strategies specifically tailored to the subject.

To address this issue, we proposed different population machine learning models to target the optimal insulin dose. The models were developed by leveraging a simulated dataset, which includes patient-specific parameters de-

scribing the mealtime condition in a noise-free environment. Both linear and nonlinear approaches were proposed, to improve the SF traditionally used for such a task and, hence, improve the quality of glycemic control while increasing model complexity. We evaluated the models' performance through an ISCT, which showed the ability of the proposed models to better approach the optimal dosage compared to the standard guidelines. Among the proposed models, we selected a final candidate model, showing the more appropriate trade-off in terms of model interpretability and performance, and we evaluated the efficacy of the selected model within a scenario which included error sources, such as the CGM measurement error or the carbohydrates counting error. The positive results obtained through the ISCT were confirmed by the aforementioned retrospective analysis on real data. Therefore, preliminary results suggested that the combination of both personalized features and machine learning-based models is key to improving the calculation of the mealtime insulin dose and further reducing the risk of hypo- and hyperglycemia. However, the design of a supervised learning framework is far from being trivial in such a context, due to the difficulty in retrieving the optimal target of the learning task in real data. Moreover, the insulin amount estimation is highly patient-dependent, making it difficult to train a population model which is tailored to different subjects.

To overcome the previous limitations, in the following study we proposed an individualized and adaptive insulin dosing technique, based on a double deep Q learning algorithm, which is tailored to the patient thanks to a personalization procedure relying on a two-step learning framework. The proposed dosing strategy was developed and tested using the UVA/Padova T1D Simulator, showing a beneficial impact on glycemic control compared to the standard guidelines.

Contents

List of abbreviations

ADA American Diabetes Association

AL Method to correct SF from Aleppo et al.

BC Bolus calculator

BG Blood glucose

BGRI Blood glucose risk index

BU Method to correct SF from Buckingham et al.

BW Body weight

CF Correction factor

CGM Continuous glucose monitoring

CHO Carbohydrates

COB Carbohydrate-on-board

CR Carbohydrate-to-insulin ratio

CSII Continuous subcutaneous insulin infusion

DSS Decision support system

DDQ Double deep Q learning

DQN Deep Q network

G_C Current glucose level

G_T Target glucose level

GBT Gradient boosted tree

HBGI High blood glucose index

IG Interstitial glucose

IB Insulin bolus

IOB Insulin-on-board

ISCT In silico clinical trial

KL Method to correct SF from Klonoff et el.

LASSO Least absolute shrinkage and selection operator

LBGI Low blood glucose index

LR Linear regression

MDI Multiple daily injections

MLR Multiple linear regression

MSE Mean-squared error

PE Method from Pettus and Edelman

RF Random Forest

ROC Rate of change

RL Reinforcement learning

SC Method to correct SF from Scheiner

SD Standard deviation

SF Standard formula

SMBG Self-monitoring of blood glucose

T1D Type 1 diabetes

TAR Time above range

TBR Time below range

TIR Time in range

ZI Method to correct SF from Ziegler et al.

List of Tables

List of Figures

Chapter 1

Insulin dosing in type 1 diabetes: standard open-loop therapy guidelines and new perspectives

1.1 Type 1 diabetes (T1D) and the standard open-loop insulin therapy

Type 1 diabetes (T1D) is a chronic disease caused by a pancreatic dysfunction, that involves progressive autoimmune destruction of the insulin-producing beta cells[1, 2]. The lack of endogenous insulin production results in excessively elevated blood glucose (BG) levels. Indeed, normal blood glucose concentrations should remain within a narrow range (70–180 mg/dl) throughout the day despite fluctuations due to nutritional intake, physical exercise, and other physiological or psychological conditions; when this concentration goes above or below the normoglycemic range, the type 1 diabetic subject experiences hyperglycemia (BG $>$ 180 mg/dL) or hypoglycemia (BG $<$ 70 mg/dL) respectively [3, 4].

In healthy individuals, optimal glycemic control is achieved by an appropriate response of insulin production from the beta cells of the pancreas, while in people living with T1D the insulin response does not occur.

An improper glycemic control may affect many different organ systems in the body, leading to serious complications over time, which can be classified as microvascular or macrovascular, and account for the major morbidity and mortality associated with T1D. Microvascular complications include retinopathy, nephropathy, and neuropathy, while macrovascular complications refer to cardiovascular, cerebrovascular, and peripheral vascular diseases [5]. Hence, to avoid such complications, patients with T1D require lifelong insulin therapy, aimed at providing for the lack of insulin in the body, to maintain their BG concentration within the normoglycemic range.

Insulin replacements include basal insulin coverage, to control the BG levels in fasting conditions, and prandial insulin boluses, to counteract the postprandial glycemic excursion. The standard guidelines for insulin therapy recommend multiple daily injections (MDIs) of basal and prandial insulin or continuous subcutaneous insulin infusion (CSII) [6, 7]. In particular, MDI therapy requires a bolus injection of long-acting insulin once or twice daily for basal insulin coverage together with short- or rapid-acting insulin at each meal, or each time a correction of high BG is required, by employing insulin syringes or pens [8].

On the other hand, CSII relies on the continuous delivery of short-acting insulin via an insulin pump, i.e. an infusion set inserted subcutaneously. All the injected boluses are composed of short-acting insulin: small amounts of insulin that are released continuously throughout the day to cover the basal insulin, while additional insulin boluses can be delivered on demand to match meal intake or to correct high BG.

Both the MDI and CSII insulin therapies are crucial to regulating BG levels

to keep subjects from getting below or above the normoglycemic range.

In particular, to counteract the glycemic excursion which follows a meal, the American Diabetes Association (ADA) provided clear guidelines regarding the meal bolus dose estimation [9], which should be computed following the formula in 1.1, hereafter defined as standard formula (SF) [10]

$$SF = \frac{CHO}{CR} + \frac{G_c - G_t}{CF} - IOB \tag{1.1}$$

where CHO (g) is the meal carbohydrate intake, CR (g/U) and CF (mg/dL/U) are the insulin-to-carbohydrates ratio and the correction factor, i.e., two therapy parameters [11], G_c (mg/dL) is the mealtime BG level, G_t (mg/dL) is the target BG level, and IOB (U) is the insulin on board [12].

The SF can be divided into three different terms. The first one, i.e., CHO/CR regards the meal insulin, which is composed of the amount of CHO consumed during the meal and the CR, which represents how many grams of CHO are covered by each unit of insulin [11]. It allows the patient to inject the insulin amount needed to compensate for the carbohydrate content planned to be eaten, e.g., a CR of 1:10 means 1 unit of insulin will cover 10 grams of carbohydrates.

The second term, i.e. (G_c - G_t)/CF, concerns the correction of insulin, it contains the difference between the current BG value and the target BG, divided by the CF, which represents how much BG is lowered by each unit of insulin [13]. Both CR and CF are patient-specific empirically estimated parameters and may vary for the same patient during the day. Furthermore, it can be noted that the meal insulin part is eliminated if there is no meal consumed by the patient and similarly the correction insulin part is deleted if the BG is already on target.

The third term consists of the IOB, indeed if insulin from previously admin-

Figure 1.1: Example of the Dexcom G6 systems, including the CGM sensor, and the auto-applicator, together with devices which allow the real-time monitoring of CGM values and trend.

istered boluses is still active, this insulin amount should be subtracted from the rest of the formula. The purpose of IOB estimations is to prevent the accumulation of multiple insulin boluses which may result in hypoglycemia.

Despite the standard guidelines provided for T1D treatment, achieving tight glycemic control could be challenging, due to the multiple factors which influence glucose regulation, e.g., hormonal changes, physical activity, and inflammatory diseases. Recent advances in T1D technologies provide the possibility of supporting the patient during the decision-making process which characterizes such a disease, thanks to the technological advances in glucose monitoring and insulin delivery, which have contributed to easing the burden related to the management of T1D [14]. Indeed, insulin delivery devices such as insulin pens and pumps, have rapidly advanced in the past two decades, to accomplish insulin delivery in a most precise manner with minimal invasiveness. On the other hand, glucose monitoring has been transformed and greatly improved by continuous glucose monitoring (CGM) devices, which automatically track BG levels throughout the day, typically with a frequency of 5-15

Figure 1.2: Representative scenarios showing the glycemic condition of the subject through the Dexcom platform, which shows in real-time the glucose levels and its ROC.

minutes [15, 16, 17], for several consecutive days. Moreover, many of the commercially available CGM systems show glucose trend arrows, i.e. CGM rate-of-change (ROC), which predict the glucose direction of travel at any given time, while providing information on the direction and magnitude of glucose change, as depicted in Figure 1.2. Having information on glucose dynamics, i.e., knowing if the BG is stable, increasing or decreasing, can help to identify and prevent unwanted periods of hypo- and hyperglycemia, which probably could not be detected with a standard self-monitoring of BG (SMBG) [18, 19].

Thus, CGM provides maximal information about shifting BG levels throughout the day and facilitates the fulfilment of optimal treatment decisions for diabetic patients. In addition, thanks to the significant refinement obtained by CGM technology [20], many devices can now be used non-adjunctively, i.e., without the need for a confirmatory SMBG for treatment decisions [21]. For this reason, exploiting the information which CGM technology provides is fundamental to optimising T1D therapy.

In particular, knowledge about glucose trends can be used to adjust the in-

sulin dose accordingly. Thus, leveraging such data could be beneficial for a more effective insulin bolus calculation and, consequently, might lead to improved postprandial glycemic control. However, the SF does not include any information on glucose dynamics at mealtime, being G_c the only term reflecting the BG status, which is a static measurement of BG concentration.

1.2 Limits and open issues of the standard guidelines for insulin therapy

Despite the technological advances in T1D management, the standard guidelines for insulin dosing present several limitations which could potentially prevent achieving optimal glucose control. First, as previously mentioned, the SF does not take into account relevant information on the glucose dynamics, provided in real-time by the CGM devices. Moreover, this formula shares the same structure among different subjects, hence not allowing for personalized insulin dosing. In addition to this lack of individualization, the standard insulin therapy greatly affects the quality of life of people living with T1D, due to the multiple tasks that are required, e.g., the meal carbohydrates counting, which may result in a patient's estimation error, leading to an improper glycemic control. Beyond this, the work of Marden et al [22] reported that a significant number of people living with T1D are unable to perform rudimental math due to limited numeracy skills.

To complicate bolus calculation even further, CR and CF parameters composing the SF may vary during the day according to intra- and interday insulin sensitivity variability, physical activity, alcohol consumption and other factors that could affect the patient metabolism, such as hormone cycles, illness, and stress [23, 24].

Figure 1.3: Examples of commercially available insulin delivery devices, starting from the standard insulin syringe, to more sophisticated systems such as the insulin pen, the jet injector and the insulin pump.

As an additional barrier to optimizing insulin therapy, the vast majority of people affected by T1D are treated with MDI therapy, which relies on injection therapy, inhaled insulin, or basic patch pumps, entailing a higher sensitivity to omitted and miscalculated doses [25]. Indeed, MDI therapy is the most common insulin therapy regimen compared to CSII and, worldwide, the usage of MDI ranges from 70% to 99% of all patients with T1D [26]. However, data collected from individuals treated under MDI therapy showed that one in four meals is associated with either a late or missed insulin bolus [27].

All the aforementioned aspects greatly impact the achievement of optimal glycemic control. Therefore, the complexity in estimating the correct amount of prandial insulin together with the multiple factors which affect such a dosage justify the use of algorithms and devices that facilitates insulin dosing, by easing the burden of this task while incorporating a more personalized and reliable calculation, which does not require any effort from the patient.

1.3 Advances in technology for T1D management

As mentioned in the previous Section, the vast majority of people living with T1D follow the MDI therapy, leaving a large population who could benefit

from the availability of tools to support daily insulin dosing but who do not want the inconvenience, cost, complexity, or commitment required with insulin pump systems. The increasing use of insulin pen technology, which supports the patient in insulin administration, fostered the investigation of smart bolus calculators to be integrated within such a device. With the popularization of smartphones, Bluetooth and near-field connectivity, a new category of smart insulin pens was released, which offer this dosing support by connecting wirelessly with a smartphone and a diabetes management app, addressing challenges in optimizing insulin injection therapy such as omitted and miscalculated doses[25] [28]. One such insulin pen is the InPen, the first Food and Drug Administration (FDA)-cleared insulin pen with a bolus calculator in the form of a smartphone app, which includes several features such as data integration with CGM devices and a detailed insulin dosing record [29, 25, 30]. The benefits associated with automated bolus calculators have been already demonstrated in multiple studies [31, 32]. Hence, developing new insulin dosing adjustment methods to overcome the limitations which characterize the standard insulin therapy is becoming a key point in T1D research.

1.4 Simulation software for the development and testing of insulin treatments

The recent development of simulation tools which mimic the glucose-insulin interaction in people affected by T1D has been of great help in advancing T1D research.

Several mathematical models have been proposed to describe the physiology of patients with T1D, including those of Hovorka et al. [33], Willinska et al. [34], Kanderian et al. [35], Haidar et al. [36], and Visentin et al. [37]. The

primary application of this simulation software consists in simulating and testing treatment protocols identical to proposed clinical studies and measuring the impact on T1D management and treatment. Indeed, such tools have promoted the investigation of new strategies aimed at improving the therapy of this disease, by allowing the simulation of realistic scenarios. For this reason, in recent years T1D research community focused on the design of novel algorithms aimed at preventing the occurrence of adverse events thanks to glucose prediction models and mitigating the entity of such events through the delivery of rescue carbohydrate intakes, corrective insulin boluses or through the optimization of mealtime insulin dose.

1.4.1 The UVA/Padova T1D Simulator

In particular, the proposed insulin dosing strategies included in this thesis were developed thanks to the use of the UVA/Padova T1D simulator [38], which relies on a comprehensive mathematical model of glucose metabolism in T1D subjects. In the following the principal components of the simulation environment that enabled the realization of the presented algorithms will be explained.

This simulator consists of 13 nonlinear differential equations, that accurately describe the metabolism of a T1D subject. These equations contain more than 30 parameters to capture the large variability in glucose dynamics among individuals with T1D (inter-patient variability). The simulator included a population of 300 in silico subjects (100 adults, 100 adolescents, and 100 children), and each virtual subject was represented by a model parameter vector, which was randomly extracted from an appropriate joint parameter distribution. Figure 1.4 depicts a scheme of the model included within the simulation software. The first version of this computer simulator is the UVA/Padova T1D Simula-

tor S2008 [39]. The updated version, the UVa/Padova T1D simulator S2013 [38], presents several additional improvements concerning S2008, both on the model on which the simulator is based, but also on the joint parameter distribution, the definition of clinically relevant parameters, and the strategy for a virtual patient generation. Its latest version, the UVa/Padova T1D simulator S2018 [37], introduces important time-varying phenomena in the simulation: a model of the so-called 'dawn phenomenon, i.e., an increase of endogenous glucose production in the early morning, and a model of intra-day variations of insulin sensitivity, which mimic the daily patterns observed in vivo. The pattern of insulin sensitivity is also subjected to random inter-day variations. In this thesis, S2013 was employed for the work presented in Chapters 3 and 4, while S2018 was for Chapter 5. Notably, in 2008 and then in its revised version in 2013, the UVA/Padova T1D model was accepted by the FDA as an alternative to pre-clinical trials.

1.4.2 Evaluating new insulin dosing strategies through in-silico clinical trials

As mentioned in the previous section, such simulation tools can be leveraged not only to develop new therapeutic strategies but also to preliminary assess their effectiveness. Indeed, testing new therapies for T1D treatment on human subjects is a time-consuming, costly and risky process. According to the Tufts Center for the Study of Drug Development, the commercialization of new pharmaceutical products has been increasing exponentially mainly because of clinical assessment [40]. Clinical assessment requires the enrollment of a sufficient number of subjects to provide proof of a new product/treatment superiority. More importantly, clinical studies have to guarantee patient safety, but this is not easy as it strongly depends on which strategy one has to test.

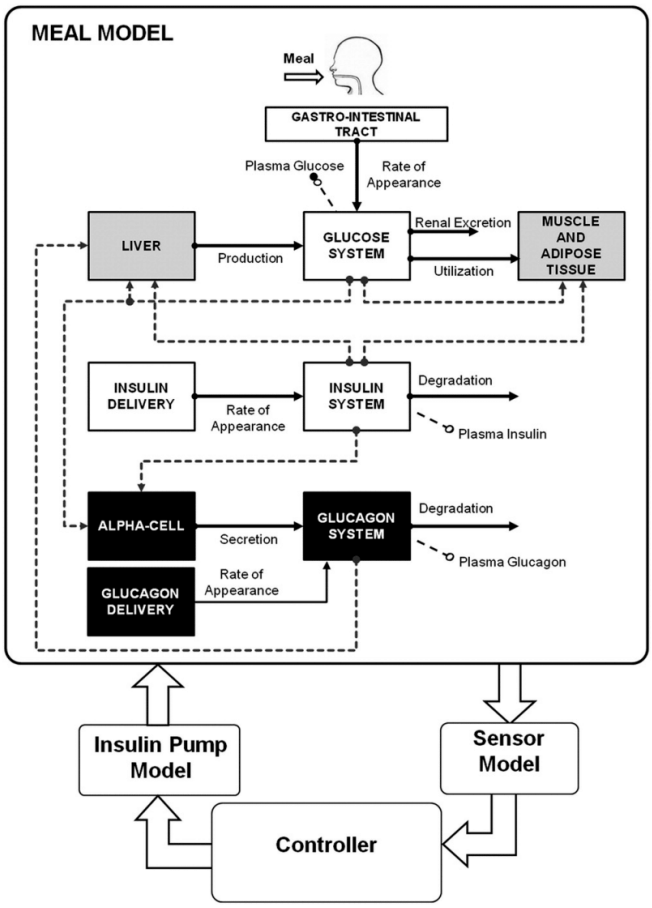

Figure 1.4: Scheme of the model included in the FDA-accepted T1DM simulator. White blocks are the unit processes of S2008 (gastro-intestinal tract, glucose kinetics and insulin kinetics); grey blocks are those that were updated in the S2013 to account for counter-regulation (liver, muscle, and adipose tissue); black blocks were added in the latest update (alpha cell, glucagon kinetics, and delivery).

To overcome these issues, in-silico clinical trials (ISCTs) can be performed. A lot of time was spent in discussing a definition of in silico clinical trials that were broad enough to cover all relevant use cases, but specific enough to be informative and useful. The consensus was reached on the following: "The use of individualized computer simulation in the development or regulatory evaluation of a medicinal product, medical device, or medical intervention. It is a subdomain of in silico medicine, the discipline that encompasses the use of individualised computer simulations in all aspects of the prevention, diagnosis, prognostic assessment, and treatment of disease" [41]. Hence, the main advantage of ISCTs is that they allow the testing of new treatments in a large number of virtual subjects, at low costs, in short times, and maintaining the same surrounding conditions.

Specifically, ISCTs could be a powerful tool to preliminary evaluate the safety and effectiveness of new insulin dosing strategies, by exploring the impact of new insulin treatments both in a low-risk situation and in numerous high-risk scenarios, which are impossible, dangerous, and unethical to replicate in a real-life setting, without any risk for real patients. For this reason, ISCTs are increasingly used to assess the performance of new therapies, medical guidelines, or algorithms for the treatment of T1D in pre-clinical trials.

1.5 Towards a personalized insulin bolus calculator

Over the last few years, the development of simulation tools has encouraged the investigation of new strategies to optimize insulin dosing. In addition, the increased amount of available data collected thanks to new technologies such as CGM sensors, insulin pens and pumps, fostered the development of smart bolus calculators integrating powerful data-driven strategies aimed at improving the standard bolus calculator based on different optimization tech-

niques. For instance, in Herrero et al. [42], run-to-run control and case-based reasoning were used to provide insulin dose recommendations based on the retrospective optimization of the therapy parameters of (1.1) (i.e., CR and CF) performed daily, while in Fabris et al. [43] the use of two methods to inform insulin dosing with biosignals from wearable sensors, i.e. insulin sensitivity estimated through CGM signal and the step count was proposed. In addition, in Cappon et al. [44], a simple neural network was preliminarily investigated for such a scope. All these studies produced positive results, encouraging further efforts on this specific topic.

As a matter of fact, these strategies have the potential to greatly improve the outcome of current standard therapy guidelines and at the same time allow to reduce the burden which characterizes the treatment of T1D, by integrating the algorithm into a smart insulin pen. Following this rationale, the design of smart bolus calculators is key to optimising glycemic control, by personalizing such a dosage both at an individual and population level. Indeed, CGM data can be used to discover patterns that can be exploited to modulate meal-insulin bolus accordingly. Furthermore, individualized parameters can be leveraged to tailor the computation of this amount to the patient's condition. Given the complexity of the problem, the nonlinear nature of glucose-insulin dynamics and the availability of a huge amount of information collected from different domains, suitable strategies, which can implement data-driven techniques to deliver personalized medical treatment, must be devised. To this end, the application of machine learning methods in T1D is presently crucial to transform the huge amount of available information into valuable knowledge.

In particular, supervised learning models, such as linear regression or decision trees, could be leveraged to target the optimal mealtime insulin dose, by using input features extracted from the CGM data, such as the ROC information, but also other easily accessible therapy parameters, such as CR, CF, or the

body weight, that describe the individual patient's physiology and characterize the glycemic response to a given amount of insulin [45].

However, the design of a supervised learning framework is far from being trivial in such a context, due to the difficulty in retrieving the optimal target of the learning task. Indeed, when considering data collected from people living with T1D, the administered mealtime insulin bolus is in most cases suboptimal, leading to poor postprandial glycemic control. Moreover, the insulin amount estimation is highly patient-dependent, making it difficult to train a general model which is valid for different subjects. Hence, the difficulty in having reliable data, which includes an optimal mealtime insulin bolus, together with the need for a model which is tailored to the patient's needs, promotes the application of reinforcement learning for such a task. Indeed, the application of reinforcement learning is particularly suitable, since it aims at automating the decision-making process characterized by T1D management, by learning through interaction with a specific environment to achieve a goal, without the need for labelled data. The previous considerations set the stage for this thesis work, which focuses on the design of smart mealtime insulin bolus calculators, by exploring both supervised learning and reinforcement learning techniques, to overcome the intrinsic limitations which characterize the standard guidelines for mealtime insulin dosing.

1.6 Aim and structure of the thesis

In this thesis, the open issues related to the standard guidelines for insulin dosing will be addressed. Particularly, this work aims to propose new strategies which tailor the insulin dosing to the patient, by applying linear and nonlinear supervised learning algorithms, as well as reinforcement learning, while integrating both information on glucose dynamics and patient-specific features.

The thesis is structured as follows: in Chapter 1, an introduction to T1D management and the related technologies is given, with a specific focus on insulin therapy. Moreover, we pointed out the key role of simulation tools in designing and testing novel insulin dosing strategies, through in silico clinical trials (ISCTs).

Then, in Chapter 2, a review of the current state-of-art methods for insulin dose adjustment was performed. Such approaches take into account the information related to the glucose dynamics, are easily accessible in real-time thanks to the use of continuous glucose monitoring (CGM) sensors and are based on empirical strategies. In addition, we proposed a novel method in collaboration with the Unit of Metabolic Diseases (Department of Medicine, Padova), which is based on the literature. The evaluation of the effectiveness of both the state-of-art methods and the proposed technique was conducted within an ISCT, employing the UVA/Padova T1D Simulator. This preliminary analysis pointed out that there is no literature method which outperforms the others in terms of postprandial glycemic control, highlighting the need for insulin dosing strategies specifically tailored to the subject and that adopt a rigorous modelling approach. The work proposed in this Chapter was previously published in [46] and [47].

Hence, in Chapters 3 and 4 we discussed the development of supervised models to improve mealtime insulin dosing, starting from the generation of the simulated dataset employed for the modelling procedure (Chapter 3) and the design and assessment of the models (Chapter 4). The dataset was created by means of the UVA/Padova T1D simulator and included patient-specific parameters describing the mealtime condition in a noise-free environment. We pointed out the importance of using a simulated dataset for such a purpose, due to the difficulty in retrieving the optimal insulin dose, i.e. the insulin amount leading to an optimal postprandial glycemic control, within records

belonging to real-world clinical data. We leveraged the generated dataset for the design of the population supervised learning models to target the optimal insulin dose. Both linear and nonlinear approaches were proposed, to improve the SF traditionally used for such a task and, hence, improve the quality of glycemic control while increasing model complexity. We evaluated the models' performance through an ISCT, which showed the ability of the proposed models to better approach the optimal dosage compared to the standard guidelines. Among the proposed models, we selected a final candidate model, showing the more appropriate trade-off in terms of model interpretability and performance, and we evaluated the efficacy of the selected model within a scenario which included error sources, such as the CGM measurement error or the carbohydrates counting error. The positive results obtained through the ISCT were confirmed by the aforementioned retrospective analysis on real data. Therefore, preliminary results suggested that the combination of both personalized features and machine learning-based models is key to improving the calculation of the mealtime insulin dose and further reducing the risk of hypo- and hyperglycemia. However, the design of a supervised learning framework is far from being trivial in such a context, due to the difficulty in retrieving the optimal target of the learning task. Moreover, the insulin amount estimation is highly patient-dependent, making it difficult to train a population model tailored to different subjects. This Chapter is based on the work published in [48, 49, 50].

To overcome the previous limitations, in Chapter 5, we proposed an individualized and adaptive insulin dosing technique, based on a double deep Q learning algorithm, which is tailored to the patient thanks to a personalization procedure relying on a two-step learning framework. The proposed dosing strategy was developed and tested using the UVA/Padova T1D Simulator, showing a beneficial impact on glycemic control compared to the standard

guidelines.

Finally, the major contributions and findings of this thesis were summarized within Chapter 6, pointing out the possible applications and future developments of the presented work.

Chapter 2

Literature approaches to adjust the standard formula (SF) for insulin dosage: review and in-silico evaluation

The advantages offered by the adoption of CGM devices in T1D therapy are remarkable since they provide not only quasi-continuous readings of glucose, but also display a trend arrow indicating its magnitude and direction, i.e., the glucose ROC. Trend arrows grant a rough short-term forecast of future glucose concentration to the user, for this reason, knowledge of trend arrows opened up the possibility of their integration within the mealtime insulin bolus (IB) calculation. As highlighted in Chapter 1, the SF does not take into account the information on glucose ROC provided by trend arrows, thus potentially leading to a suboptimal dosage. In general, a positive ROC may suggest that the IB dose should be increased, while, on the other hand, a negative ROC indicates that the dose should be reduced. Even if intuitive, providing clear and effective recommendations on how to adjust the IB based on ROC is far

from trivial, since under/over dosages could potentially lead to suboptimal glycemic control and, in some cases, to critical glycemic levels [51]. Hence, the need for precise guidelines together with the availability of trend arrows fostered the development of several methodologies aimed at adjusting the SF by considering the ROC. However, a comprehensive comparison of the performance and safety of such methods is still missing. Designing a trial to answer this question could not be easy, since comparing several methods for IB calculation under the same mealtime conditions could be practically impossible. This problem can be circumvented by resorting to ISCTs, to draw preliminary indications. An ISCT for such a purpose was designed by Cappon et al. [52], where the UVA/Padova T1D Simulator [38] was used to test mealtime insulin dosing strategies accounting for ROC in the same scenario. However, [52] the evaluation was limited to three literature methodologies available at that time, while in the last year several other methods were published [53, 54, 55, 56, 57]. Hence, the study proposed in this Chapter aims to perform a more extensive comparison, including, in addition to the three methods originally considered, other four recently published methods, reviewed in the following Section 2.1. Moreover, a novel adjustment method will be proposed in Section 2.2 developed in collaboration with the Unit of Metabolic Diseases at the Department of Medicine (University of Padova). The literature methods, together with the proposed one, will be evaluated in Section 2.3 thanks to an ISCT, using the UVa/Padova T1D simulator in a single-meal scenario. Lastly, results in terms of safety obtained from the analysis of real data, where our method was employed from the patients, will be reported in Section 2.4.

The work presented in this Chapter was published in the papers of Noaro et al. [46] and Bruttomesso et al. [47].

2.1 Review of the current literature methods for insulin dose adjustment

We considered the three methods by Buckingham et al. (BU) [57], Scheiner (SC) [53], Pettus/Edelmann et al. (PE) [56], already considered in [52], and the three recent contributions by Klonoff/Kerr (KL) [58], Aleppo/Laffel et al. (AL) [55], and Ziegler et al. (ZI) [54]. The analysed literature methods correct the insulin dose based on the glucose ROC, sharing the same division into five separate intervals, which are based on the direction and magnitude of glucose change. Particularly, if the ROC takes values between -1 and +1 mg/dL, it is considered stable, and no correction is needed, while adjustments are applied for positive (i.e., [+1, +2] mg/dL/min and [+2, +3] mg/dL/min) and negative (i.e., [-1, -2] mg/dL/min and [-2, -3] mg/dL/min) ROC values. To summarize the methodologies, we classified them into three categories based on the different approaches adopted to adjust SF according to ROC.

2.1.1 Method based on a percentage modulation of SF

This category contains only BU, which is the first published guideline for meal-time IB adjustment using ROC. The authors suggested adjusting the SF of Eq. 1.1 by applying a percent modulation proportional to the ROC value. Of note, it has been shown that such modulation is perceived as too modest from the patient perspective, who usually prefers larger adjustments [59].

2.1.2 Methods based on the adjustment of current glucose in SF

This category, which includes SC and PE methods, exploits the notion of anticipated glucose, i.e., the predicted glucose value in 30-60 minutes given G_c,

and ROC. This interval approximately corresponds to the time required by rapid-acting insulin analogue to affect the glucose concentration, in addition, the 30-60 minutes timeframe is short enough to assume that the glucose trend will be stable within that interval. Thus, SC and PE methods followed this rationale to adjust the G_c used within the SF 1.1according to ROC magnitude and direction by increasing/decreasing its value. Particularly, the SC approach is more conservative compared to PE, since the former proposes adjustments lower in module compared to the latter.

2.1.3 Methods that correct SF by a fixed amount

The three previous methods could be burdensome for T1D individuals, especially for those who lack numeracy skills and may experience difficulties in estimating the right dose due to the required calculations [22]. For this reason, KL, AL, and ZI work proposed a simplified approach, which consists in modifying the SF by a fixed insulin amount, both without considering personalized information of the T1D individual, as in KL, or adjusting also based on a personalized therapy parameter, i.e., CF, as in AL, and ZI.

We refer the reader to Tables 2.1 and 2.2 for more details on these methodologies.

2.2 Design of a novel empirical method (Bruttomesso et al.)

The insulin adjustment method proposed in the work of Bruttomesso et al. [47] (BR) is based on the literature correction previously proposed by Ziegler et al. [54], but narrower intervals for starting BG values and CF were considered. In particular, the proposed method relies on the same division in ROC intervals,

while modifying the CF subdomains, by considering five distinct classes of insulin sensitivity factors: less than 30, between [30 - 40], between [40 - 60], between [60 - 90] and higher than 90 mg/dL/U. In addition, compared to Ziegler et al., the prandial glucose value was divided into three possible classes, i.e., [70 - 120], [120 - 180], [180 - 250] mg/dL. Narrower intervals were considered for both CF and G_c to better tailor the insulin correction to the patient's conditions. Further details on the proposed method are reported in Table 2.3, which shows the correction value for each interval.

Table 2.1: Guidelines for IB adjustment related to the works of BU, SC, PE, KL, and AL. The reported values are limited to the cases of interest (ROC $>$ 3 mg/dL/min or $<$ -3 mg/dL/min are excluded

		Adjustment method							
		BU	SC	PE	KL	AL			
	CF					<25	25-50	50-75	>75
Positive ROC	2-3	+20% of IBSF	+50 mg/dL to GC	+75 mg/dL to GC	+1.5 U	+3.5 U	+2.5 U	+1.5 U	+1 U
[mg/dL/min]	1-2	+10% of IBSF	+25 mg/dL to GC	+50 mg/dL to GC	+ 1 U	+2.5 U	+1.5 U	+1 U	+0.5 U
Stable ROC [mg/dL/min]	<1	0	0	0	0	0			
Negative ROC	1-2	-10% of IBSF	-25 mg/dL to GC	-50 mg/dL to GC	-1 U	-2.5 U	-1.5 U	-1 U	-0.5 U
[mg/dL/min]	2-3	-20% of IBSF	-50 mg/dL to GC	-75 mg/dL to GC	- 1.5 U	-3.5 U	-2.5 U	-1.5 U	-1 U

Table 2.2: Guidelines for IB adjustment related to the work of ZI. The reported values are limited to the cases of interest (ROC > 3 mg/dL/min or < -3 mg/dL/min are excluded, as well as prandial BG > 250 mg/dL).

	Glucose level								
		70-180 mg/dl				180-250 mg/dl			
	CF	<25	25-50	50-75	>75	<25	25-50	50-75	>75
Positive ROC	2-3	+2.5	+ 2	+ 1	+ 0.5	+3.5	+ 2.5	+ 1.5	+ 1
[mg/dL/min]	1-2	+1.5	+ 1	+ 0.5	+ 0.5	+2.5	+ 1.5	+ 1	+ 0.5
Stable ROC [mg/dL/min]	<1	0	0	0	0	0	0	0	0
Negative ROC	1-2	–2.5	–1.5	–1	–0.5	–2	–1	–0.5	–0.5
[mg/dL/min]	2-3	–3.5	–2.5	–1.5	–1	–3	–2	–1	–1

Table 2.3: Guidelines for IB adjustment related to the work of BR. The reported values are limited to the cases of interest (ROC > 3 mg/dL/min or < -3 mg/dL/min are excluded, as well as CF > 90 and prandial BG > 250 mg/dL).

	Glucose level												
		70-119 mg/dl				120-180 mg/dl				181-250 mg/dl			
	CF	<30	30-40	40-60	60-90	<30	30-40	40-60	60-90	<30	30-40	40-60	60-90
Positive ROC	2-3	+2	+ 1.5	+ 1	+ 0.5	+2.5	+ 2	+ 1.5	+ 1	+ 3	+ 2.5	+ 2	+ 1.5
[mg/dL/min]	1-2	+1	+ 1	+ 0.5	+ 0.5	+1.5	+ 1.5	+ 1	+ 1	+ 2	+ 2	+1.5	+ 1.5
Stable ROC [mg/dL/min]	<1	0	0	0	0	0	0	0	0	0	0	0	0
Negative ROC	1-2	–2	–2	–1.5	–1.5	–1.5	–1.5	–1	–1	–1	–1	–0.5	–0.5
[mg/dL/min]	2-3	–3	–2.5	–2	–1.5	–2.5	–2	–1.5	–1	–2	–1.5	–1	–0.5

2.3 Evaluation of methods' effectiveness through an in-silico clinical trial

2.3.1 Simulation environment

Each methodology was assessed through ISCTs in a simulated environment, being such a framework suitable for this type of analysis, where a virtual cohort of T1D individuals underwent different IB adjustments maintaining the same scenario. The UVA/Padova T1D Simulator (T1DS) [38] was used, and the virtual cohort included 100 adult subjects, which assumed a range of CF values from 26 to 67 mg/dL/U. Briefly, T1DS [60], consists of a mathematical model describing the glucose-insulin dynamics in people with T1D using 13 differential equations and more than 30 parameters to represent the large inter-subject variability. T1DS is equipped with 100 virtual subjects, each associated with a different realization of the parameter set, which can be used to assess the performance of new strategies for T1D management by designing ad-hoc ISCTs. The software has been accepted in 2008 by the Food and Drug Administration as an alternative to preclinical trials and it is widely used by the diabetes technology research community [48, 52, 43, 61, 62]. Within this framework, each subject underwent multiple single-meal ISCTs, lasting 12 hours, from 7 am to 7 pm. The first timeframe (from 7 am to 1 pm) was exploited to bring the subject to specific prandial conditions. Particularly, we simulated different scenarios in terms of ROC, ranging between -2 and +2 mg/dL/min with a step of 0.5 mg/dL, and BG, taking values of 80, 120, 160, and 200 mg/dL. We did not cover ROC values higher than 2 mg/dL and lower than -2 mg/dL, since those values were not easily obtainable through realistic actions (e.g., small CHO intakes or insulin boluses) assumed in the preprandial window. Then, a meal was set at 1 pm, when each virtual subject had a carbohydrate

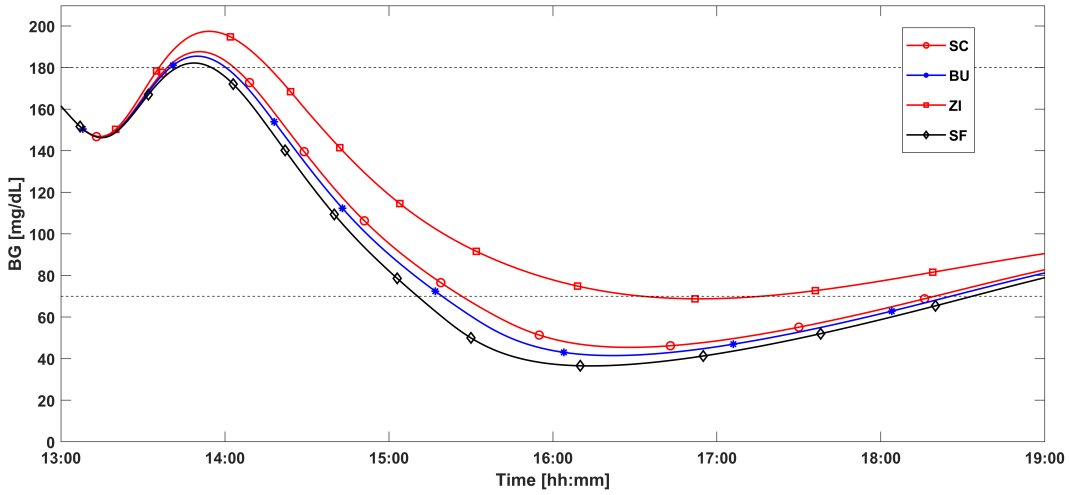

Figure 2.1: Representative examples of BG curves during postprandial time window for different methods of insulin bolus computation and different mealtime conditions. Mealtime ROC is negative (-1.5 mg/dL/min), starting BG=160 mg/dL and meal CHO is 60 g. The calculated IB doses are IB_{SF}=3.62 U, IB_{BU}=3.26 U , IB_{SC}=3.02 U, IB_{ZI}=2.12 U.

intake composed of different amounts (from 10 to 150 g, with a step of 10 g) and the corresponding IB, computed using the methodologies under assessment (SF, BU, SC, PE, KL, AL, ZI, BR), was tested for each prandial condition. The simulation lasted for a postprandial interval of 6 hours (from 1 pm to 7 pm), in which glucose fluctuations were not affected by any corrective action. Moreover, within the experimental set-up, we did not consider any source of error, i.e., BG measurement error, ROC estimation error, CHO counting error, nor variability, i.e., insulin sensitivity, to evaluate only the contribution given by the literature methods. Thus, for each prandial status, which is defined by a specific combination of ROC and BG at mealtime, 1500 glycemic traces were generated, resulting from 15 different CHO amounts for every virtual subject. A representative example of the simulated mealtime scenario is reported in Figure 2.1, where, for better visualization, only SF, SC, BU and ZI insulin dosages were applied.

2.3.2 Evaluation metrics and statistical analysis

For the sake of simplicity, results were grouped into two different main scenarios based on prandial ROC values, i.e., negative (-2, -1.5, -1 mg/dL) and positive (1, 1.5, 2 mg/dL), to assess the benefit of a decreased and increased IB dose separately. Moreover, we evaluated the literature methods performances within the 6-hour postprandial interval of each simulation, by computing standard metrics that quantify glucose control, such as the BG risk index (BGRI) [63], the percentage of time spent within the target glycemic range (TIR), i.e., BG between 70-180 mg/dl, above the range (TAR), i.e., BG > 180 mg/dl and below the range (TBR), i.e., BG < 70 mg/dl [64, 65]. To better highlight the possible improvement concerning SF, we calculated the point differences between each metric obtained with the literature methods and SF (ΔBGRI, ΔTIR, ΔTAR, ΔTBR). In addition, summary results of every single metric distribution will be presented in the median and interquartile range. The statistical significance was evaluated on the single metric distributions, by applying Friedman's test with a 5% significance level. We used this nonparametric test, due to the non-Gaussian metric distributions and the repetition of the subjects within the dataset. Moreover, the p-values resulting from the statistical test were adjusted using the Bonferroni correction to account for multiple comparisons.

2.3.3 Results

Differences between the metric distributions (ΔBGRI, ΔTIR, ΔTAR, ΔTBR) are shown in Figures 2.2, 2.3, 2.4, and 2.5, for the two scenarios, i.e., positive and negative ROC, and for each prandial BG value considered in the study, i.e., 80 mg/dL, 120 mg/dL/ 160 mg/dL, 200 mg/dL. In the figures, we highlighted with green/red backgrounds the regions in which the literature method led to an improvement/worsening of glucose control versus SF respectively. More-

over, in Tables 2.4 and 2.5 the resulting median and interquartile ranges of the single metric distributions for each method are reported. For each scenario, the method leading to the best glucose control was selected and highlighted with bold text within Tables 2.4 and 2.5. The selection was performed by looking first at those minimizing BGRI, which is a global metric considering both the risk of hyperglycemia and hypoglycemia, in presence of similar BGRI values, also TAR, TIR and TBR were taken into account in the selection process.

Negative ROC scenario

As shown on the left side of Figures 2.2, 2.3, 2.4, and 2.5, similar glycemic control was obtained when the ROC is negative for all considered metrics and all BG values. In particular, it was generally found that ΔTAR < 0 (red area), indicating an increased TAR compared to SF. On the other hand, ΔTBR was mostly above 0 (green area), showing an improvement in TBR for all methods versus SF. This result was expected, since a negative ROC drives to a lower IB amount compared to SF, promoting the shortcoming of hyperglycemic episodes. Moreover, ΔTIR and ΔBGRI improved for all the BG values compared to SF. The overall improvement of the latter metric can be explained by the greater risk associated with hypoglycemia concerning hyperglycemia within the BGRI. Finally, it can be noticed that the more the starting BG is higher, the more the improvement in terms of ΔBGRI, ΔTBR, and ΔTIR is evident. Analyzing the results of the single metrics reported in Table 2.4, the following considerations can be made: BG = 80 mg/dL: All the methods produced lower TBR, and an improved TIR compared to SF. However, despite the such improvement, no statistical difference was detected between the metric distributions obtained with the literature methodologies and SF. The method of KL achieved the highest median TIR (60.66% compared to 55.40% of SF) and the lowest BGRI (from 9.62 to 8.94). BG = 120 mg/dL: Methods of PE, KL, AL, ZI, and BR obtained

higher TIR compared to SC and BU, while all the approaches reached a median value of TBR equal to 0%, with AL and ZI having a significant reduction compared to SF. The methods leading to the lowest BGRI values were BR and PE, with BR reaching the highest median TIR (65.93%). BG = 160 mg/dL: Also, in this case, the most moderate improvement was given by BU and SC, while the outcomes obtained by PE, KL, AL, ZI, and BR are more pronounced, especially in terms of median TBR, reporting median values of 0%, which are significantly lower compared to SF. The worsening in TAR showed a significant difference from SF for KL and ZI. The method providing the best performance in terms of TIR and BGRI, without significantly increasing the TAR, resulted in PE. BG = 200 mg/dL: The benefits provided by the correction of SF are more evident, indeed the BGRI values obtained by PE, KL, AL, and ZI are significantly lower than those of SF, as well as the improvement of TIR and TBR. Methods of AL and PE achieved a median TBR equal to 0%, reaching the lowest BGRI values and the highest TIR (53.19% and 50.97% respectively). We observed, however, a moderate increase of TAR, which became significant only for AL, suggesting an under-correction by such a method. For this reason, PE led to the best performance, since it did not significantly increase the TAR.

Table 2.4: Quantitative assessment of glycemic control when prandial ROC is negative. Median and interquartile ranges of BGRI, TAR, TIR, and TBR are reported for each state-of-art method and SF, according to the prandial value of BG (80, 120, 160, 200 mg/dL). Bold text indicates the best performing methods within the prandial ROC and BG subdomain.
*statistically significant compared to SF.

		Negative ROC			
BG [mg/dL]		BGRI	TAR	TIR	TBR
	BU	8.99 [4.64-16.47]	13.57 [0-30.33]	59.83 [38.23-81.99]	7.76 [2.77-39.89]

	SC	8.97 [4.64-16.66]	15.79 [0-30.75]	59.83 [37.67-81.44]	6.93 [2.49-40.44]
	PE	9.06 [4.82-15.99]	19.39 [0-34.63]	60.39 [39.89-78.95]	4.99 [1.94-35.18]
	KL	**8.94 [4.79-16.18]**	**18.56 [0-33.52]**	**60.66 [39.06-79.22]**	**5.26 [1.94-37.53]**
	AL	9.24 [4.93-16.34]	20.22 [0-36.01]	59.56 [39.06-78.12]	4.99 [1.66-34.9]
	ZI	9.24 [4.93-16.34]	20.22 [0-36.01]	59.56 [39.06-78.12]	4.99 [1.66-34.9]
	BR	9.47 [5.12-16.38]	22.16 [0-37.12]	58.73 [39.61-76.45]	4.71 [1.39-32.55]
	SF	9.62 [4.75-18.14]	12.19 [0-28.25]	55.4 [34.9-81.44]	18.56 [3.6-44.04]
120	BU	8.35 [3.81-15.18]	20.78 [0-31.86]	63.16 [42.11-86.43]	0 [0-32.41]
	SC	8.09 [3.63-15.54]	21.61 [0-31.58]	63.99 [41-87.53]	0 [0-32.69]
	PE	7.79 [3.83-13.93]	25.48 [0-35.18]	65.65 [47.37-85.32]	0 [0-21.33]
	KL	7.86 [3.81-14.23]	24.65 [0-34.35]	65.65 [45.71-85.87]	0 [0-24.93]
	AL	8.03 [3.88-13.95]	26.59 [0-36.57]	65.1 [47.37-83.93]	0* [0-17.59]
	ZI	8.03 [3.88-13.95]	26.59 [0-36.57]	65.1 [47.37-83.93]	0* [0-17.59]

	BR	**7.85**	**25.21**	**65.93**	**0**
		[3.85-14.01]	**[0-34.9]**	**[47.09-85.32]**	**[0-22.16]**
	SF	9.34	18.84	57.06	19.53
		[3.97-18.28]	[0-29.64]	[36.84-86.15]	[0-38.78]
160	BU	9.28	25.21	59.97	5.26
		[4.31-16.59]	[0-34.07]	[37.95-84.49]	[0-32.41]
	SC	9.09	25.21	59.28	6.65
		[3.92-17.79]	[0-32.96]	[35.73-87.26]	[0-34.07]
	PE	**8**	**28.53**	**66.07**	**0***
		[3.9-14.79]	**[1.94-36.01]**	**[43.07-84.76]**	**[0-22.99]**
	KL	8.18	27.7	64.82	0*
		[3.87-15.25]	[0-35.46]	[41-86.01]	[0-26.32]
	AL	8.15	29.64*	65.37	0*
		[4.06-14.42]	[6.65-37.67]	[45.71-83.38]	[0-18.01]
	ZI	8.15	29.64*	65.37	0*
		[4.06-14.42]	[6.65-37.67]	[45.71-83.38]	[0-18.01]
	BR	8.03	28.25	65.93	0*
		[3.89-14.82]	[0-36.01]	[42.94-85.04]	[0-23.55]
	SF	11.58	22.99	46.54	28.95
		[4.91-21.91]	[0-31.02]	[32.13-81.72]	[0-40.72]
200	BU	13.19	34.07	37.67	26.32
		[6.89-21.92]	[22.44-41.27]	[25.76-65.37]	[0-37.67]
	SC	13.71	33.24	34.9	28.81
		[6.44-24.75]	[22.71-39.89]	[23.55-67.87]	[0-39.89]
	PE	**10.69***	**36.01**	**50.97***	**0***
		[5.31-19.34]	**[26.59-42.66]**	**[28.12-70.91]**	**[0-31.58]**

KL	11.17* [5.51-20.38]	35.46 [25.48-42.11]	46.81* [26.87-70.64]	8.73* [0-34.07]
AL	10.38* [5.46-18.35]	37.12* [27.98-44.04]	53.19* [29.92-69.53]	0* [0-28.67]
ZI	11.6* [5.57-20.98]	35.18 [25.21-41.83]	44.88* [26.04-70.64]	12.47* [0-34.9]
BR	12.56 [5.88-22.85]	34.07 [23.82-40.72]	38.23 [24.65-70.36]	23.27 [0-37.67]
SF	17.73 [8.79-30.3]	31.58 [20.5-38.23]	29.36 [21.05-52.35]	37.4 [18.28-45.71]

Positive ROC scenario

By observing the right side of Figures 2.2, 2.3, 2.4, and 2.5, similar, specular, considerations to the previous scenario can be made for all BG values while considering positive ROC values. As expected, since in such a scenario all methods led to a higher IB dosage, the ΔTAR improved (green area), while the ΔTBR generally showed positive distributions (red area), indicating an increased number of hypoglycemic episodes induced by the considered methods concerning SF. The medians ΔBGRI and ΔTIR resulted mostly above and below zero (red areas), respectively, suggesting a general worsening of the overall glycemic control, especially when high mealtime BG values were considered. By analyzing Table 2.5, the following considerations can be made: BG = 80 mg/dL: The TIR of all the literature methods did not differ significantly from SF, as well as the BGRI distributions. The 75th percentile of TBR increased, maintaining the median to 0%, on the contrary, the TAR decreased for all the methods. The increase in TBR was found significant only for PE and AL, likewise the reduction in TAR. The method having the best performance

in terms of TIR and BGRI proved to be BR, despite the moderate improvement (TIR from 61.22% of SF to 63.16%, BGRI from 9.61 of SF to 8.86). BG = 120 mg/dL: Methods of PE, KL, AL, ZI, and BR were shown to be overly aggressive, by significantly increasing the median TBR. Despite the TAR improvement for each method, all the TIR distributions led to a lower median value than the one of SF. The two methods that maintained a median value comparable to the one of SF (59%) are the most conservative ones, i.e., BU and SC, with the latter reaching the highest value (57.62%). BG = 160 mg/dL. Methods of PE, KL, Al, ZI, and BR induced a significant worsening of TBR compared to SF. Moreover, AL resulted particularly aggressive, by significantly decreasing the TIR (from 51.25% of SF to 31.3%). The most conservative methods, i.e., BU and SC, provided a small improvement in TAR (from 38.23% of SF to 35.46 and 35.73%, respectively). On the other hand, being more conservative, allowed BU and SC not to significantly increase the TBR. BG = 200 mg/dL: The BGRI, TIR and TBR distributions considerably worsened compared to SF for PE, KL, AL, ZI, and BR, suggesting an overcorrection of the IB amount. Only the TAR improved for all the methods, taking values between 29.09-31.58% from 33.24% of SF. Also, in this case, BU and SC proved to be the safest methods, with the latter reaching the best trade-off when considering all the metrics.

Table 2.5: Quantitative assessment of glycemic control when prandial ROC is positive. Median and interquartile ranges of BGRI, TAR, TIR, and TBR are reported for each state-of-art method and SF, according to the prandial value of BG (80, 120, 160, 200 mg/dL). Bold text indicates the best performing methods within the prandial ROC and BG subdomain.
*statistically significant compared to SF.

		Positive ROC			
BG [mg/dL]		BGRI	TAR	TIR	TBR
	BU	9.33 [5.22-15.79]	32.41 [24.1-40.17]	60.94 [41.83-75.35]	0 [0-13.85]

	SC	9.02 [4.74-15.26]	31.86 [21.61-40.44]	62.88 [45.71-77.56]	0 [0-4.71]
	PE	9.17 [4.27-17.32]	28.53* [16.34-36.29]	61.22 [37.67-81.72]	0* [0-28.25]
	KL	9.21 [4.46-16.96]	29.09 [17.31-36.84]	61.5 [38.78-79.92]	0 [0-26.59]
	AL	10.08 [4.57-18.85]	27.42* [14.4-34.9]	55.4 [34.9-80.33]	0* [0-33.24]
	ZI	9.05 [4.37-16.58]	29.36 [17.73-37.12]	62.05 [39.61-80.06]	0 [0-24.65]
	BR	**8.86 [4.49-15.48]**	**30.75 [19.94-38.78]**	**63.16 [43.49-78.95]**	**0 [0-16.9]**
	SF	9.61 [5.46-14.93]	35.18 [26.04-44.04]	61.22 [48.48-73.68]	0 [0-0]
120	BU	10.83 [5.94-19.9]	32.41 [26.04-39.34]	55.96 [33.24-72.58]	0 [0-29.09]
	SC	**10.62 [5.6-18.37]**	**32.41 [24.65-40.17]**	**57.62 [36.01-73.68]**	**0 [0-25.21]**
	PE	12.5 [5.91-22.43]	29.64* [21.05-36.84]	44.6 [30.47-73.68]	19.67* [0-36.29]
	KL	11.92 [5.72-21.35]	30.19 [21.88-37.4]	47.92 [31.58-74.24]	11.63* [0-34.63]
	AL	13.73 [6.54-24.36]	28.81* [19.94-36.01]	41 [29.09-70.36]	26.32* [0-39.61]
	ZI	11.69 [5.68-20.91]	30.47 [22.16-37.67]	49.72 [31.86-74.24]	7.48 [0-33.52]

	BR	12.27 [5.79-22.11]	29.64* [21.05-36.84]	45.71 [31.02-73.96]	18.01* [0-35.73]
	SF	10.49 [6.2-17.04]	35.46 [28.25-43.21]	59 [42.11-70.91]	0 [0-8.17]
160	BU	14.53 [8-25.89]	35.46 [29.92-41.83]	40.44 [26.04-64.54]	14.54 [0-36.57]
	SC	**14.13 [8.02-23.25]**	**35.73 [29.92-42.94]**	**43.77 [28.25-64.82]**	**8.03 [0-33.24]**
	PE	17.46 [9.28-28.58]	33.24 [27.15-39.89]	33.24 [24.38-57.34]	30.19* [0-41.55]
	KL	16.56 [8.97-27.29]	33.8 [27.7-40.44]	34.63 [25.21-60.39]	27.42* [0-40.17]
	AL	19.01* [10.26-30.81]	32.69* [26.18-39.34]	31.3* [23.55-51.52]	34.07* [0-43.21]
	ZI	16.04 [8.68-26.83]	34.07 [27.98-41]	35.46 [25.48-61.22]	26.04* [0-39.06]
	BR	17.12 [9.23-28.27]	33.52 [27.15-40.17]	33.8 [24.65-58.45]	29.36* [0-41.55]
	SF	13.18 [8.03-20.82]	38.23 [32.41-45.43]	51.25 [32.41-64.82]	0 [0-24.38]
200	BU	22.54 [11.23-39.39]	30.75 [22.71-36.84]	24.93 [18.01-39.61]	42.38 [29.92-49.86]
	SC	**21.12 [11.61-35.43]**	**31.58 [22.71-37.95]**	**25.76 [19.39-39.34]**	**41.27 [29.92-48.48]**
	PE	27.24* [16.4-44.39]	29.92 [20.5-36.29]	22.71* [17.45-31.86]	46.81* [39.34-53.46]

Table 2.6: Summary of the obtained results. For each BG and ROC scenario, the adjustment method which led to the best outcome within the ISCTs is reported

	BG = 80	BG = 120	BG = 160	BG = 200
Negative ROC	KL	BR	PE	PE
Positive ROC	BR	SC	SC	SC

KL	26.03* [15.21-42.17]	30.33 [21.05-36.57]	23.27* [17.73-32.96]	45.71* [37.67-52.63]
AL	29.65* [18.17-46.44]	29.36 [19.94-35.73]	22.16* [17.17-30.47]	48.48* [40.72-54.85]
ZI	29.65* [18.17-46.44]	29.36 [19.94-35.73]	22.16* [17.17-30.47]	48.48* [40.72-54.85]
BR	31.01* [19.44-48.5]	29.09* [19.39-35.18]	21.61* [16.62-29.36]	49.31* [42.11-55.96]
SF	16.96 [8.43-28.92]	33.24 [24.65-39.61]	29.36 [21.05-54.29]	35.18 [8.86-44.32]

In conclusion, Table 2.6 reports a summary of the results obtained with the simulations, which allows better identify the most effective correction method for each prandial BG and ROC scenario we tested.

2.4 Safety assessment of Bruttomesso et al. method using real data

Following the ISCT conducted within the UVA/Padova T1D Simulator, a preliminary assessment to evaluate the safety of the Bruttomsso et al. method was performed. In particular, data collected from 27 adult subjects, that underwent periodic medical examinations in the medical centres involved in the study, were analyzed retrospectively. The subjects followed the insulin therapy

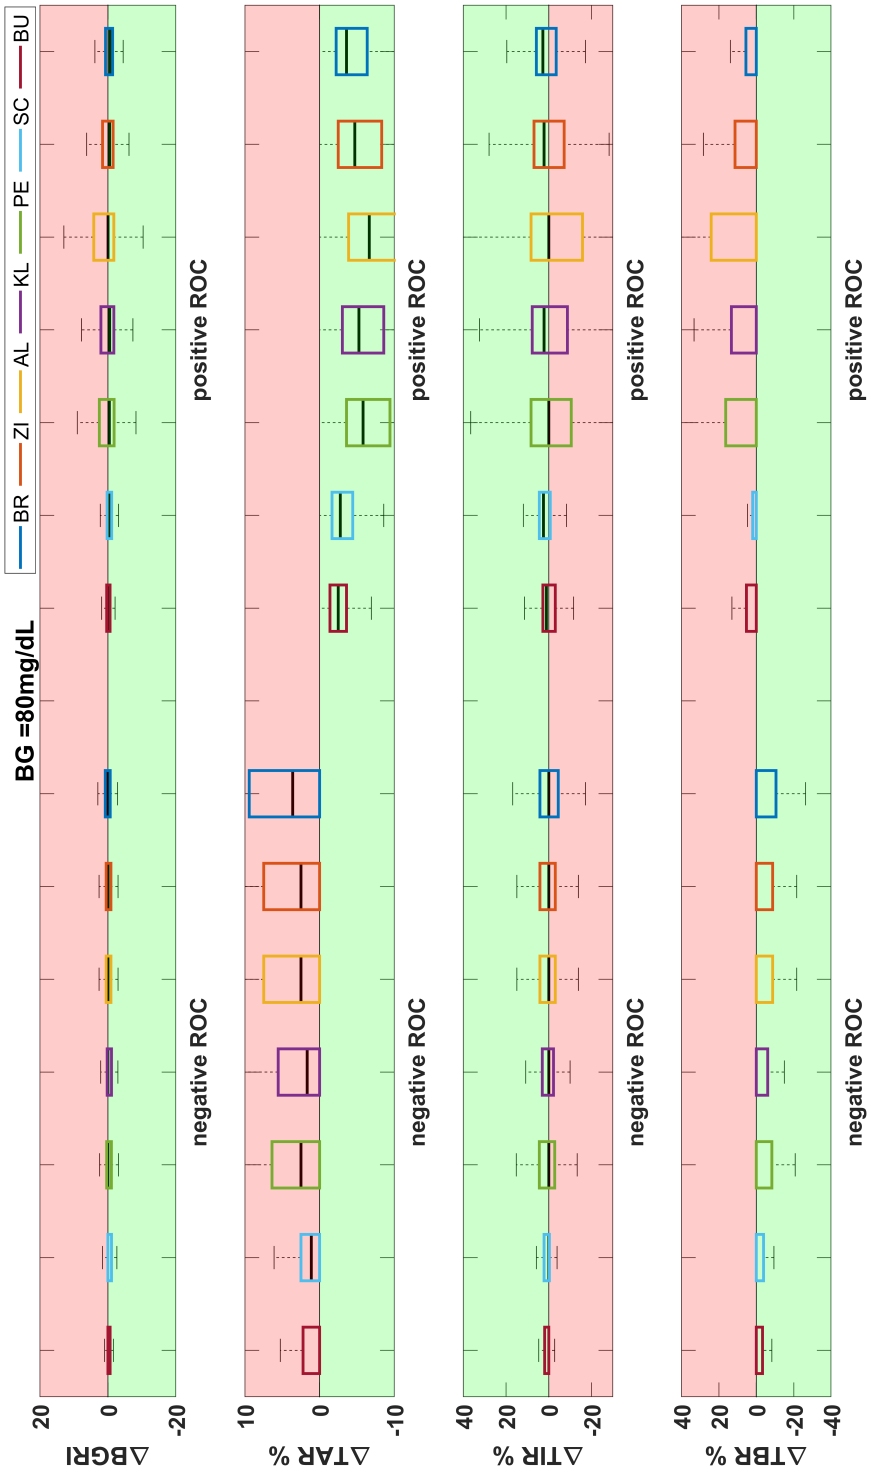

Figure 2.2: Distribution of ΔBGRI, ΔTAR, ΔTIR, ΔTBR (difference between the literature methods and SF) for negative (left) and positive (right) ROC with a prandial BG of 80 mg/dL. The green background corresponds to an improvement of the method concerning SF, while the red background corresponds to a worsening.

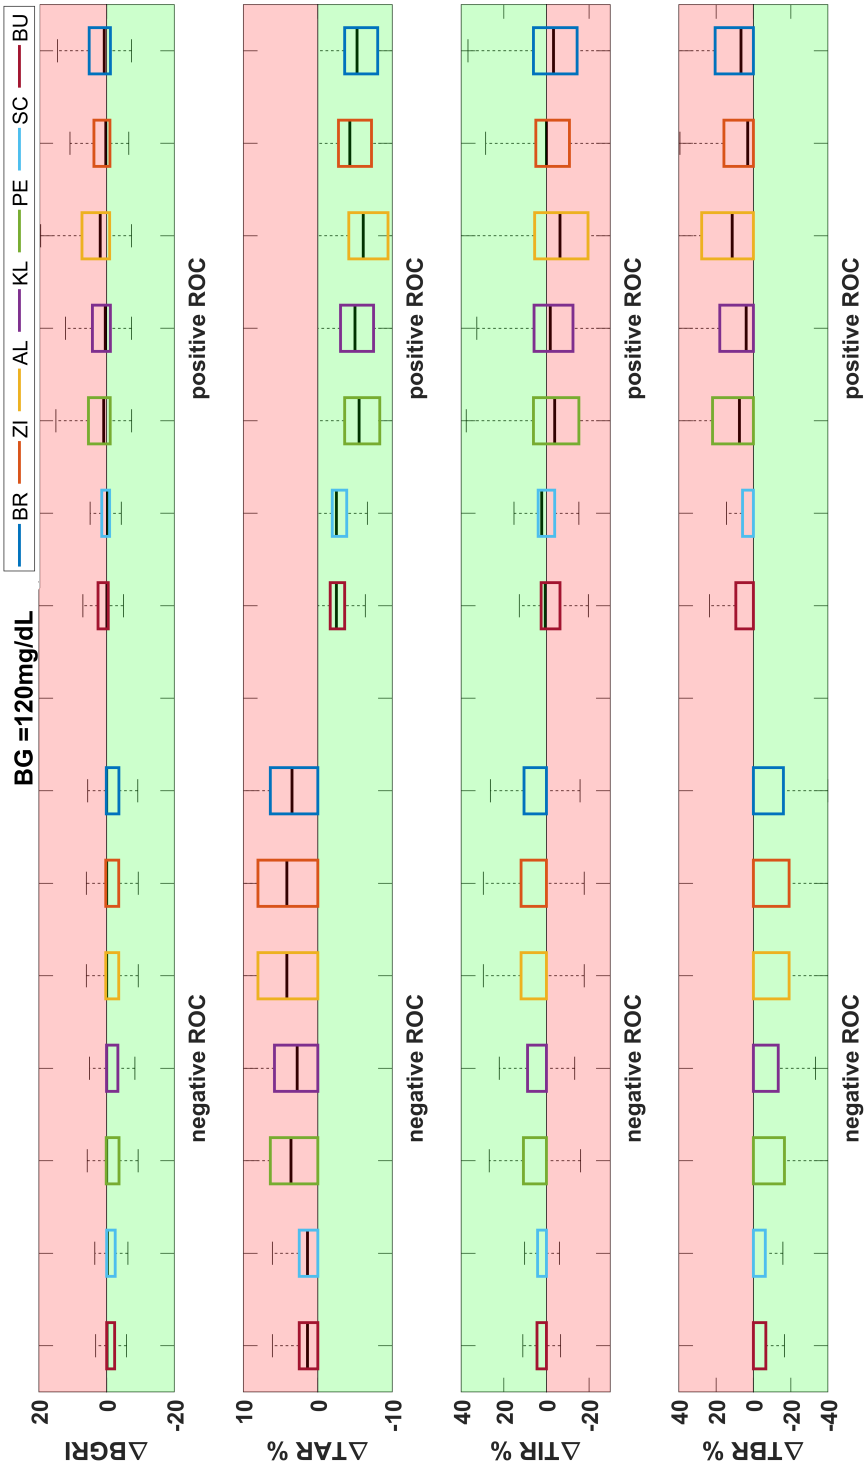

Figure 2.3: Distribution of ΔBGRI, ΔTAR, ΔTIR, ΔTBR (difference between the literature methods and SF) for negative (left) and positive (right) ROC with a prandial BG of 120 mg/dL. The green background corresponds to an improvement of the method concerning SF, while the red background corresponds to a worsening.

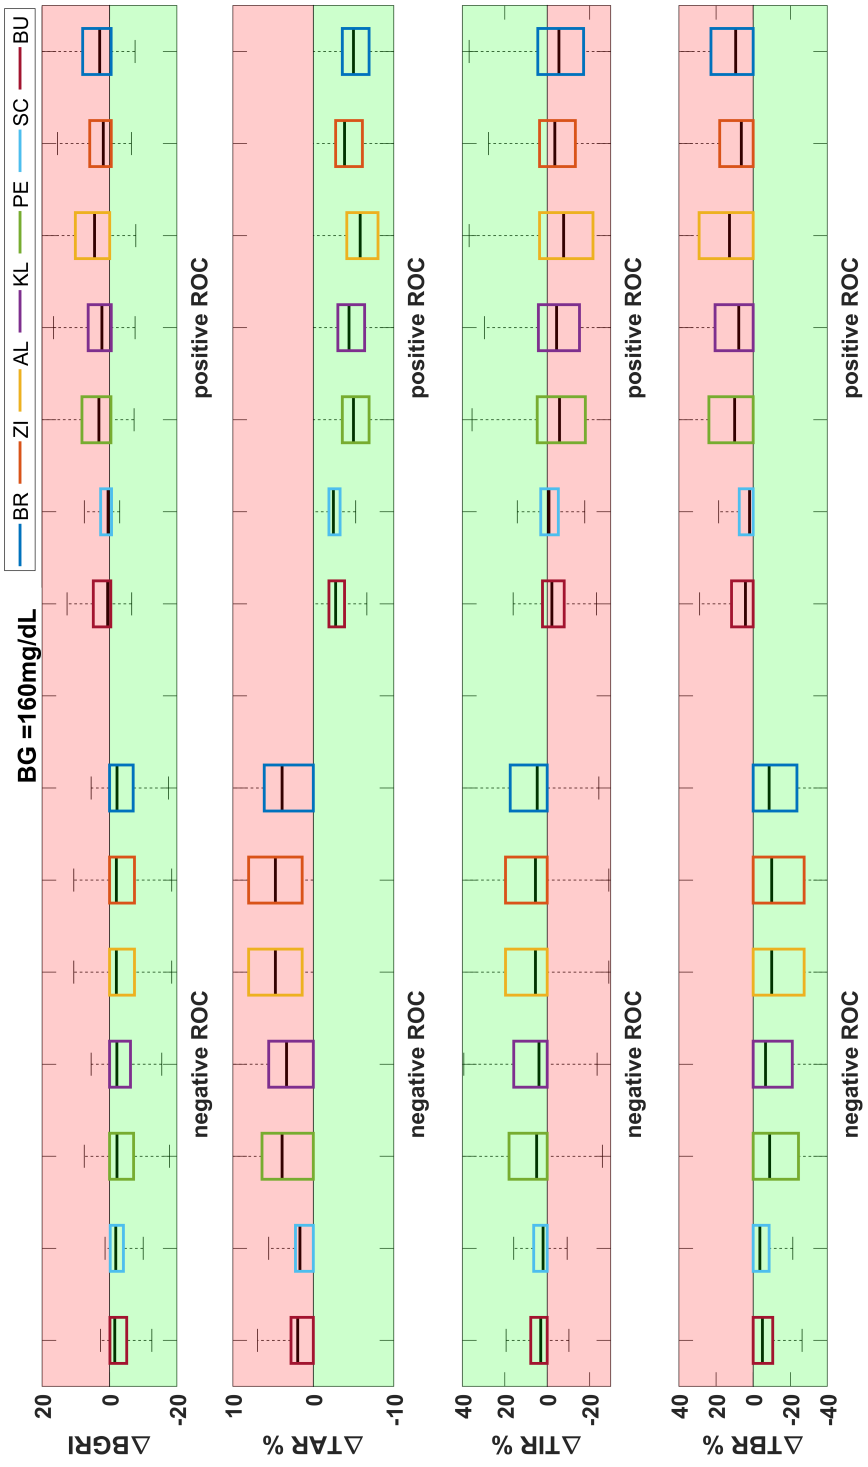

Figure 2.4: Distribution of ΔBGRI, ΔTAR, ΔTIR, ΔTBR (difference between the literature methods and SF) for negative (left) and positive (right) ROC with a prandial BG of 160 mg/dL. The green background corresponds to an improvement of the method concerning SF, while the red background corresponds to a worsening.

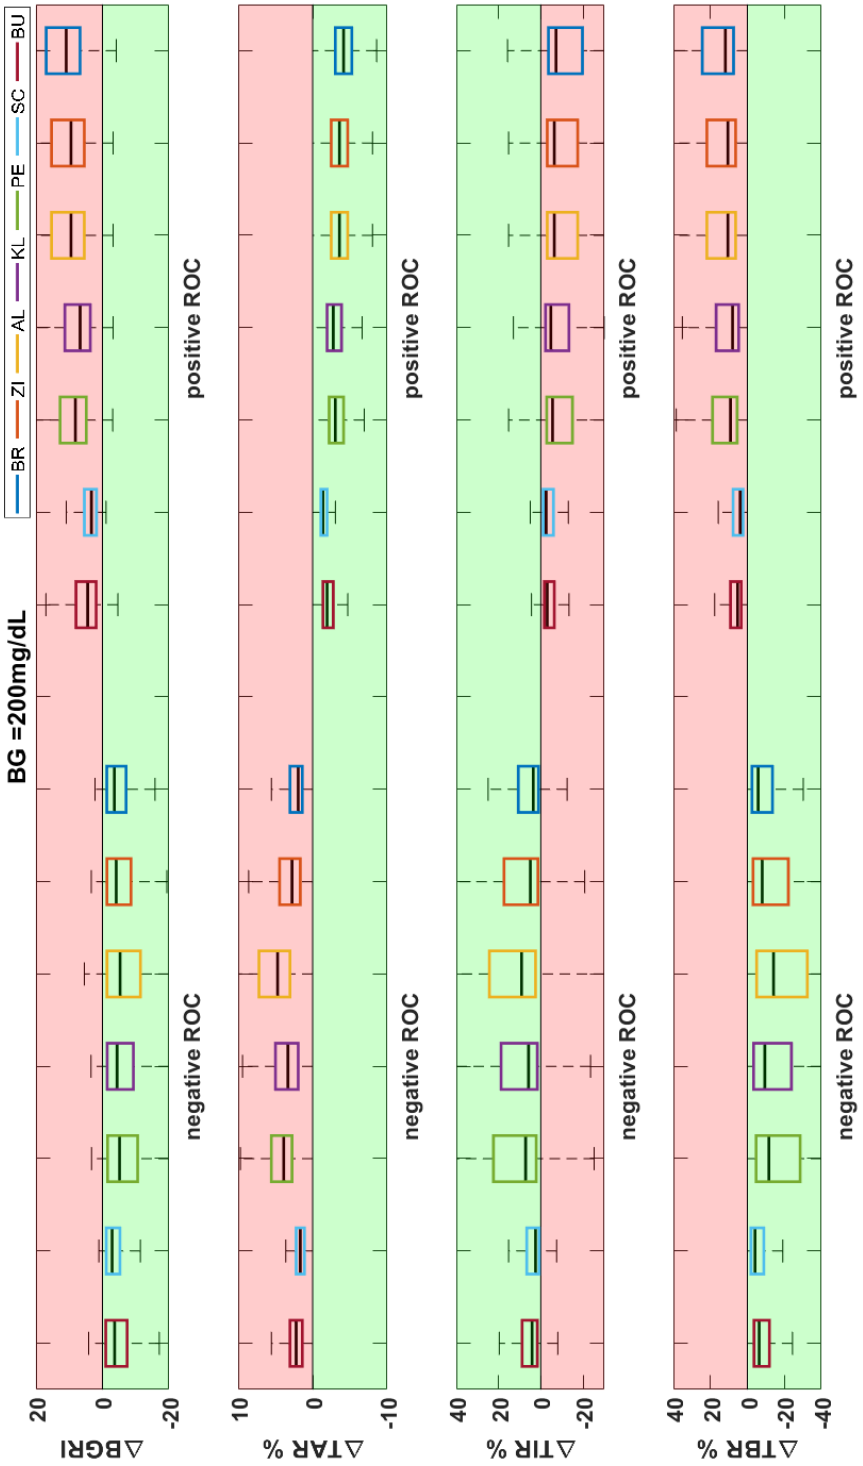

Figure 2.5: Distribution of ΔBGRI, ΔTAR, ΔTIR, ΔTBR (difference between the literature methods and SF) for negative (left) and positive (right) ROC with a prandial BG of 200 mg/dL. The green background corresponds to an improvement of the method concerning SF, while the red background corresponds to a worsening.

suggested by the proposed method for at least 2 weeks and kept track of information regarding the CHO amount of the meal, prandial CGM, SMBG, trend arrow and insulin doses. Among the 27 patients, 16 were male, mean (± SD) age was 49.3 ± 13.3 years, mean duration of diabetes was 27.7 ± 13.3 years, and mean glycated haemoglobin (HbA1c) was 55.7 ± 6.9 mmol/mol (7.2 ± 2.8%). Nine and eighteen patients were on MDI and CSII therapy, respectively. The proposed insulin therapy was followed for a mean (± SD) of 13.4 ± 6.4 days. We evaluated only the meals in which the standard insulin bolus was estimated correctly, the ROC was above 1 mg/dl/min or below - 1 mg/dl/min, and no corrective actions were present during the 4-hour postprandial time window.

For each extracted CGM profile, TAR, TBR, and TIR were computed in the 4-hour postprandial time window. We reported the results in two scenarios: (1) increasing ROC only and (2) decreasing ROC only. Finally, for each scenario, we computed TAR, TBR, and TIR considering three different time intervals: (1) from mealtime to 4 h after the meal; (2) from mealtime to 2 h after the meal; and (3) from 2 to 4 h after the meal, hereafter referred to as T0–4, T0–2, and T2–4 , respectively. Analysis of T0–4 helps understand if the meal bolus was accurate. To evaluate the bolus effect early or late after the meal, we considered the 0–2 h and 2–4 h postprandial intervals, respectively. Data are reported as the median and corresponding interquartile ranges.

2.4.1 Results

Among all the extracted meals, only 172 were preceded by an increasing or decreasing arrow. Of these, 64 were not considered in the analysis because the standard insulin bolus had been calculated erroneously or meal composition had not been recorded. Consequently, the final dataset comprised data from

108 meals, 52 with increasing and 56 with decreasing ROC. In 85% of cases, the increasing ROC ranged between 1 and 2 mg/dl/min, 13% ranged between 2 and 3 mg/dl/min, and the remaining 2 % had values less than 3 mg/dl/min. Regarding the decreasing ROC, 93 % of the values ranged between 1 and 2 mg/dl/min, while 7 % ranged from 2 to 3 mg/dl/min. For the 5 % of meals initiated with sensor glucose readings of 70 mg/dl, 66 % of the patients started with CGM in the euglycemic range, and 27 % started with CGM between 180 and 262 mg/dl. When the preprandial CGM was 70 mg/dl, the patients performed an SMBG. In all of these cases, the glucose levels turned out to be 70 mg/dl, and the patients calculated the insulin bolus according to the slide rule, based on the SMBG readings. The median values, together with the interquartile ranges of TAR, TIR, and TBR are reported in Table 2.7. In the increasing ROC scenario, TAR was 45.8 (0–77.0) during T0–2, decreasing to 0.0 (0.0–56.3) during T2–4. We obtained a TIR of 54.1 (22.9–100.0) during T0–2 that reached 83.3 (43.7–100.0) during T2–4. The median TBR and its interquartile ranges were zero for all the time intervals analyzed. Regarding the decreasing ROC scenario, TIR was 97.9 (72.9–100) during T0–2 while the median values of TBR and TAR were both 0. During T2-4, the median TBR and TAR were also both equal to 0, reaching a median TIR of 100. In general, considering the T0–4 interval in both the increasing and decreasing ROC scenarios, the median TIR was 70.8 and 91.6, respectively.

2.5 Summary of the obtained results and limitations

The analysis performed in this Chapter pointed out that there is no literature method that is globally the most effective. However, by investigating the results grouped by BG and ROC subdomains, we noticed that some methods are more effective and safer than others. In general, when negative ROCs are

Table 2.7: Results obtained from the real life use of slide rule in the two scenarios: (i) increasing ROC only, and (ii) decreasing ROC only. Median [interquartile ranges] are shown for TAR, TBR, and TIR evaluated in T0-4, T0-2, T2-4.

	Metric	T0-4	T0-2	T2-4
Increasing ROC	TAR %	27.08 [8.33-61.46]	45.83 [0-77.08]	0 [0-56.25]
	TIR %	70.83 [38.54-88.54]	54.17 [22.92-100]	83.33 [43.75-100]
	TBR %	0 [0-0]	0 [0-0]	0 [0-0]
Decreasing ROC	TAR %	1.04 [0-26.04]	0 [0-20.83]	0 [0-43.75]
	TIR %	91.67 [69.79-100]	97.92 [72.92-100]	100 [56.25-100]
	TBR %	0 [0-0]	0 [0-0]	0 [0-0]

considered, the resulting reduction of IB dosage suggested by all the methodologies proved to be beneficial in terms of glucose control, increasing BGRI and TIR, and decreasing TBR, with a modest increase in TAR. In this scenario, BU and SC were found to be systematically too conservative, leading to a minor improvement compared to the other methods. In contrast, the benefits provided by PE, KL, AL, ZI, and BR are more evident. We selected KL as the best performing method for a low starting BG value (80 mg/dL), BR for a BG value of 120 mg/dL, and PE for both BG approaching the hyperglycemic range (160 mg/dL) and BG in hyperglycemia (200 mg/dL). On the other hand, when the prandial ROC is positive, our results showed the potential risk introduced by the increase of the IB dose. In general, TIR, BGRI, and TBR worsened concerning SF. For a low prandial BG (80 mg/dL) BR resulted in the best performing recommendation, while for higher starting BG values, the most conservative methods SC and BU proved to be safer.

In our opinion, the results presented in this Chapter should be used more qualitatively than quantitatively, knowing that ad-hoc clinical trials to further validate the effectiveness of the methods are required. However, the indica-

tions are clear and solid, suggesting that, in general, decreasing IB for negative trend arrows is safe, while increasing IB when the trend arrow is positive could not be. The analysis showed that there is no method that is the best performing in all scenarios, and allowed identifying, each prandial glucose and trend arrow, which are the methods that could provide more benefits in terms of glucose control. Therefore, being the ROC adjustment a function of prandial glucose and trend arrow, a hybrid solution such that proposed in Table 3, which combines the best-performing methods for each prandial condition, is the one that in our opinion should be suggested. From a practical perspective, asking the T1D individual to apply different rules based on prandial conditions can be not straightforward. However, the adoption of a hybrid method can be made easy e.g., using mobile apps able to get real-time data from CGM sensors and automatically provide to the user the correct adjusted dose, without requiring any user intervention. In conclusion, the analysis conducted represents a step forward to close the gap present in the literature, by providing more information about the practical use of methods to adjust IB accounting for trend arrows, thus helping to define clear and safe guidelines for people with T1D for insulin dosing adjustments.

However, the derivation of all previous rules for SF correction has mainly been empirical, suggesting that there would be room for improvement should a systematic modelling methodology be adopted. Moreover, the results presented in this Chapter, pointed out the need for more effective, and possibly personalized, IB calculation strategies. Some recent proof-of-concept studies have also shown that machine learning techniques can be used to tailor the SF correction to meet patient-specific parameters [44]. Further improvement could be achieved by designing strategies which do not merely correct the SF, but define completely new rules or models for insulin dosing. Such novel approaches could leverage the information available on glucose dynamics, pro-

vided by CGM devices, but also other easily accessible patient-dependent variables, such as body weight (BW) and insulin basal rate (I_b).

Chapter 3

Leveraging a simulated dataset to target the optimal insulin dose

In this Chapter, the generation of the simulated dataset employed for the model's training will be presented. In particular, we discuss the procedure that allowed us both to simulate different mealtime scenarios, from which features describing the prandial status are extracted and derive the optimal mealtime insulin dose. Then, a correlation analysis is performed to investigate the association among the feature variables and with the optimal insulin amount. The resulting dataset will be used within the following Chapter for the training of the linear and nonlinear supervised learning models to target the optimal prandial insulin amount.

3.1 Rationale

Despite many efforts that have been made in the literature to design an empirical approach for standard insulin dose correction, the use of supervised learning techniques for such a task remains poorly investigated. The main reason lies in the impossibility of retrieving the optimal insulin bolus, i.e. the insulin

amount leading to an optimal postprandial glycemic control, within records belonging to real datasets. Indeed, although the availability of data collected from people affected by T1D, such as insulin data and CGM samples, has been recently increasing, the mealtime insulin dose is mainly suboptimal, making it unreliable to be leveraged as the target of the learning task. Hence, the advent of simulation tools which mimic the glucose-insulin interaction enabled us to generate virtual scenarios in which the optimal prandial insulin dose associated with a specific mealtime condition can be computed retrospectively, by optimizing the postprandial glucose control. Indeed, developing new mealtime insulin amount models within a simulation environment is particularly advantageous for two main reasons. Firstly, a simulation environment makes it possible to generate a unique dataset where patients undergo multiple meal tests while maintaining the same surrounding conditions. This would be impossible to replicate with clinical trials since a patient's behaviour and physiological state do not remain the same. Secondly, extreme conditions, which are difficult and dangerous to obtain in clinical trials, can be simulated without any risk to the patient. In the following, we discuss the generation of the simulated dataset, composed of features representing the mealtime status and the corresponding optimal insulin bolus associated with the specific condition.

3.2 Generation of the simulated dataset

The UVA/Padova T1D Simulator has been used to generate data from 100 adult virtual subjects. This virtual population underwent multiple single-meal scenarios in a noise-free setting [60], which consisted of: using optimal therapy parameters; not permitting either postprandial correction boluses or rescue carbohydrate intakes; and, having no errors in either CHO counting, BG measurements or ROC estimation. Moreover, we did not consider the intra-patient

variability of insulin sensitivity during the meal. All these choices were made to eliminate any confounding factors that could have influenced the outcomes of the study. Considering all these factors, we generated for each virtual subject a simulated scenario lasting 12 hours, i.e. from 7 am to 7 pm.

3.2.1 Generation of the meal conditions

The goal of this simulation is to obtain a static dataset of different conditions of the subjects, about the values of ROC and BG at a given time (1 pm). We fixed the values of BG and ROC that each subject should reach at 1 pm:

- BG = [70, 80, 90, 100, 110, 120, 130, 140, 150, 160, 170, 180] mg/dL
- ROC = [-2, -1.5, -1, -0.5, 0, 0.5, 1, 1.5, 2] mg/dL/min

The BG values cover all the euglycemic ranges with steps equal to 10 mg/dL, while the ROC arrows range from -2 to +2 with a step of 0.5 mg/dL/min.

Given the BG and ROC conditions that each subject should have at 1 pm, we needed to manipulate the meals (and the respective boluses) and snacks (that are intended as small meals with no subsequent boluses) within the previous 6-hour long time window (from 7 am to 1 pm). The number of possible meals consumed during that time window was fixed to 1 and equally was performed for the number of snacks. We defined two variables characterizing the possible preprandial meal and snack: the time when the meal and the snack are eaten from the virtual subject and the amount of CHO of which the meal and the snack are composed. Then, we determined other two variables describing the mealtime insulin bolus, i.e., the time when the bolus is injected and the quantity of injected insulin.

At this point, the values to be found to reach the aforementioned conditions are 6: time of the meal and grams of meal's CHO (t_{CHO_1}, CHO_1 respectively);

time of the meal-bolus and quantity of injected insulin (t_B, B); time of the snack and grams of snack's CHO (t_{CHO_2}, CHO_2). To derive these values, we formulated the problem as a minimization procedure, thus minimizing the following function:

$$\hat{p} = \underset{p}{\operatorname{argmin}} f(t_{CHO_1}, t_{CHO_2}, t_B, CHO_1, CHO_2, B) \tag{3.1}$$

with

$$f(t_{CHO_1}, t_{CHO_2}, t_B, CHO_1, CHO_2, B) = (BG_c - BG(p))^2 + (ROC_c - ROC(p))^2$$

where p is the vector of the 6 variables, BG_c and ROC_c are the reference values of BG and ROC set before. 3.1 is a function of the 6 variables that provide as output the exact values which allow the single subject to reach the specific BG and ROC condition.

Hence, by adopting this strategy, we simulated for each adult several mealtime scenarios, in which at 1 pm the subject experienced the ROC and BG values defined before, by acting on the meal, the bolus and the snack.

3.2.2 Generation of the optimal bolus dose

In the following, the strategy implemented to obtain the targets (optimal boluses) for each mealtime condition is explained. After the procedure performed in the previous section, we fixed the values of CHO that compose the meal consumed by each subject at 1 pm:

- CHO = [0, 10, 20, 30, 40, 50, 60, 70, 80, 90, 100, 110, 120, 130, 140, 150] g

Then, given the amount of CHO eaten at lunch, we computed the optimal insulin bolus that should be administered to the subject after the meal. The insulin bolus is considered optimal because it is calculated in a way that minimizes the Blood Glucose Risk Index (BGRI).

The BGRI, first introduced by Kovatchev et al. [66], is composed of the sum of two terms: the Low Blood Glucose Index (LBGI) and the High Blood Glucose Index (HBGI). The definition of LBGI and HBGI derives from a logarithmic transformation of the BG scale that balances the amplitude of hypo- and hyperglycemic ranges (enlarging the former and shrinking the latter) and makes the transformed data symmetric around zero and fitting a normal distribution [67]. This symmetrization is performed because in the standard BG scale, hypoglycemia (BG < 3.9 mmol/l) and hyperglycemia (BG > 10 mmol/l) have very different ranges, and euglycemia is not central in the entire blood glucose range (1.1-33.3 mmol/l). Consequently, the scale is not symmetric and its clinical centre (6-7 mmol/l) is far from its numerical centre (17 mmol/l). As a result, a logarithmic data transformation that matches the clinical and numerical centre of the BG scale has been applied, thus making the transformed data symmetric [66]. If BG measurements are expressed in mg/dl, the transformed data ($f(BG)$) are given by:

$$f(BG) = 1.509 \cdot ([log(BG)]^{1.084} - 5.381) \tag{3.2}$$

and are used to define a BG risk function:

$$r(BG) = 10 \cdot f(BG)^2 \tag{3.3}$$

that associates to each BG reading a measure of its risk, as expressed with a number in the 0 to 100 range. For the definition of this metric, the risk index of the equation 3.3 is used and thus the $LBGI$ and $HBGI$ are computed and added to generate the overall $BGRI$:

$$LBGI = \frac{1}{n} \sum_{i=1}^{n} [r(BG(i))][f(BG(i)) < 0] \tag{3.4}$$

$$HBGI = \frac{1}{n}\sum_{i=1}^{n}[r(BG(i))][f(BG(i)) > 0] \tag{3.5}$$

$$BGRI = LBGI + HBGI \tag{3.6}$$

The BGRI has been chosen as a function to be minimized since it balances hypo- and hyperglycemic events, despite their original scale are not symmetric. Through the minimization of the $BGRI$ function, we derived the target needed to reach the goal of this thesis, each variable of this dataset will be explained in detail in the following section.

3.3 Features description and dataset structure

In the previous section, the method to obtain the simulated scenarios describing the mealtime condition was discussed, while in this section all the variables extracted, that were selected as part of the dataset, will be described. The employed dataset is composed of 11 variables, the first five parameters are patient-specific, hence being constant for each virtual patient, while the last six features describe the mealtime status:

- CR: patient-specific therapy parameter, representing how many grams of CHO are covered by each unit of insulin.
- CF: patient-specific therapy parameter, representing how much the BG is lowered by each unit of insulin.
- BW: easily accessible variable indicating the patient's body weight.
- Ib: feature describing the basal insulin rate of the subject.
- Gt: feature representing the basal glucose value of the subject.
- BG: feature representing the BG level at mealtime, which can take 12 values ranging between 70 and 180 mg/dL.

- ROC: feature describing the glucose trend at mealtime, which can take 9 values ranging between -2 and +2 mg/dL/min.
- CHO: feature representing the amount of consumed CHO at mealtime, which can take 16 values ranging between 0 and 150 grams.
- SF: variable describing the amount of mealtime insulin suggested by the standard therapy, computed as in 1.1.
- COB: variable indicating the carbohydrates-on-board (COB) at mealtime.
- IOB: variable indicating the amount of insulin still acting in the organism at mealtime.

Thus, each record composing the resulting dataset was associated with the corresponding optimal insulin bolus that should be administered to the subject following the fixed lunch amount. By considering the different combinations among the possible CHO, ROC and BG variables, there are 1728 possible conditions for each patient. Hence, the total number of records which compose the dataset is 172800, i.e. 1728 possible mealtime statuses for each subject. In the following section, the relationship between the main features and the target will be visualized.

3.4 Dataset visualization

To investigate and better understand the dataset structure, we visualized different mealtime conditions which are included within the resulting database. In Figure 3.1 the BG trace of subject 64, having prandial BG and ROC equal to 100 mg/dl and 1 mg/dl/min is shown. In this specific case, the meal amount corresponds to 50 g of CHO, and the optimal bolus dose (3.1157 U), previously extracted from the optimization procedure, is injected at mealtime (1 pm). In

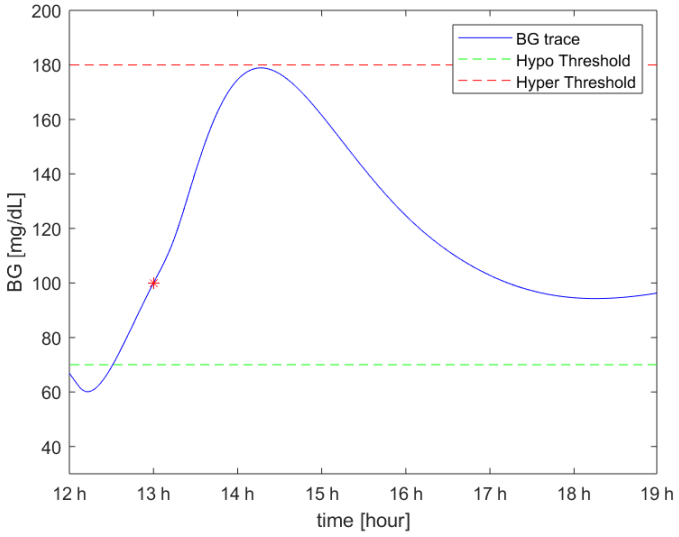

Figure 3.1: Representative BG curve of subject 64, having the following mealtime conditions at 1 pm: BG=100 mg/dl, ROC=1 mg/dl/min, CHO=50 g, B=3.1157 U/min.

Figure 3.1, the aforementioned prandial conditions are depicted. In particular, it can be seen that the BG curve reached the specified value at 1 pm, having an increasing trend. Thanks to the optimal injected bolus, no hypo- or hyperglycemia events occurred to the subject. A different scenario is shown in Figure 3.2, where the prandial BG is equal to 120 mg/dl and the ROC is falling with a value of -1 mg/dl/min. A meal composed of 50 g of CHO is consumed by the subject and the optimal injected bolus allows to maintain the BG trace within the normoglycemic range. In Figure 3.3 and 3.4 the relationship between optimal bolus, CHO values and BG is investigated. The visualization of the optimal bolus distribution was done by fixing the ROC value, to analyze different conditions of the BG trend separately. This analysis has been carried out to investigate the role of the current BG level in the amount of optimal injected insulin.

Figure 3.3 shows the data corresponding to a negative ROC value equal to -2 mg/dL/min. When the glucose trend is falling with this slope, the optimal bolus distribution has a median equal to 0 until 40 grams of CHO, while from

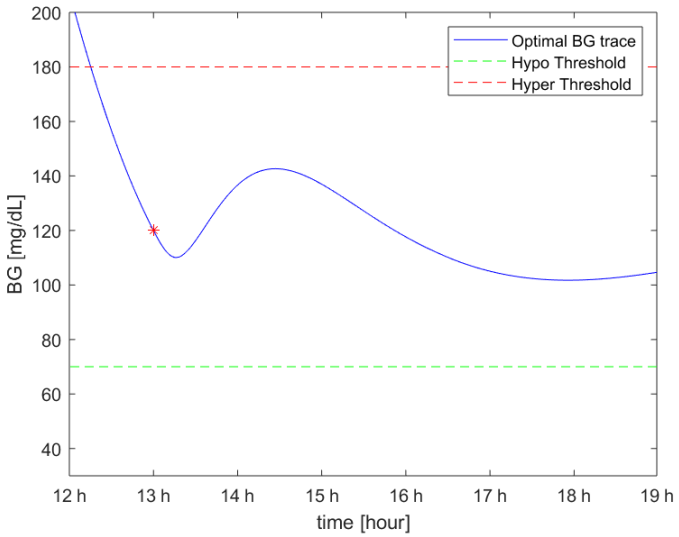

Figure 3.2: Representative BG curve of subject 64, having the following mealtime conditions at 1 pm: BG=120 mg/dl, ROC=-1 mg/dl/min, CHO=50 g, B=1.4734 U/min.

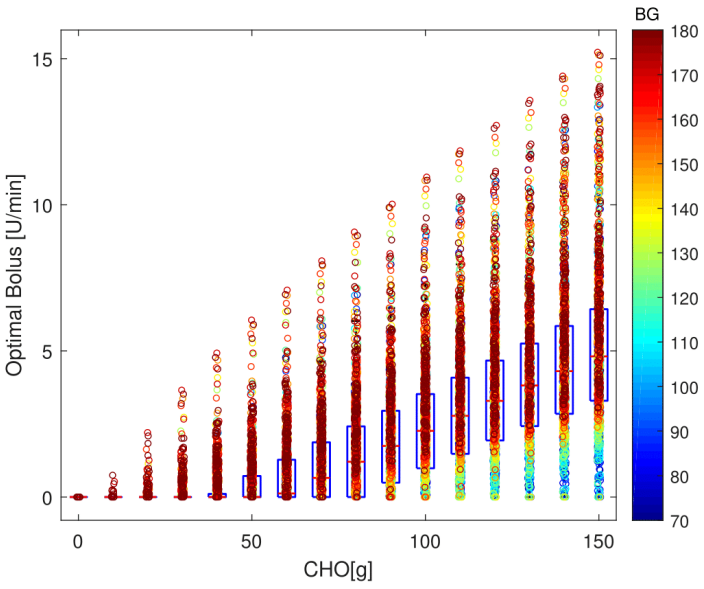

Figure 3.3: Distribution of the optimal bolus depending on carbohydrate values with ROC fixed to -2 mg/dl/min. The optimal boluses are reported for all the subjects, overlapped to the boxplot, and coloured according to the current BG value.

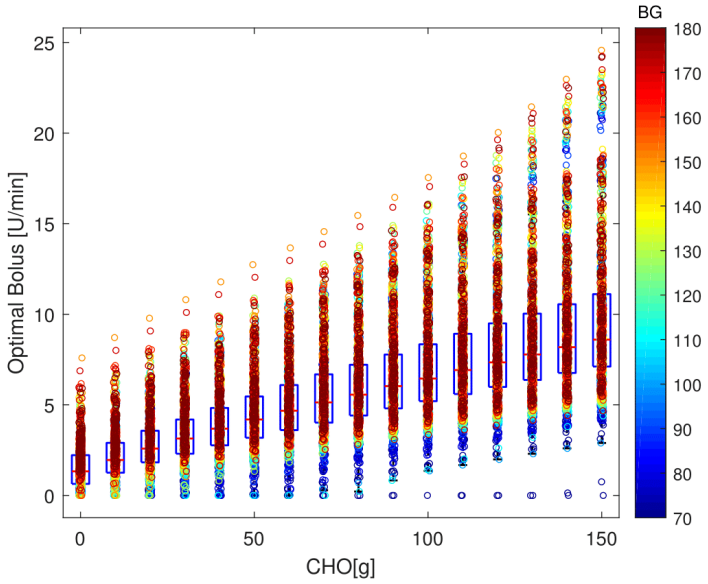

Figure 3.4: Distribution of the optimal bolus depending on carbohydrate values with ROC fixed to +2 mg/dl/min. The optimal boluses are reported for all the subjects, overlapped to the boxplot, and coloured according to the current BG value.

50 to 150 grams the median has increasing values with a linear trend, reaching 4.8 U/min for 150 gr. When the subject has such a glucose trend and there is no consumed meal, the optimal bolus is equal to 0, independently of the BG value. For small amounts of carbohydrates (from 10 to 40 g), an insulin bolus should be administered only to the subjects who have a high BG level (150-180 mg/dl). It can be noted a clear distinction between boluses with a high BG level and low BG level, especially for a significant amount of consumed CHO (100-150 g).

When a positive ROC occurs (in particular, ROC = 2 mg/dl/min), the median values of the optimal bolus distribution are shifted upwards, showing also in this condition a linear trend (Figure 3.4). Consequently, the subject needs an insulin dose even though the amount of CHO is very small (or null). Compared to the previous case (ROC = -2 mg/dL/min), the clear distinction between boluses with low BG and high BG starts with lower CHO values. It can be stated that the amount of CHO and BG values highly affect the meal-

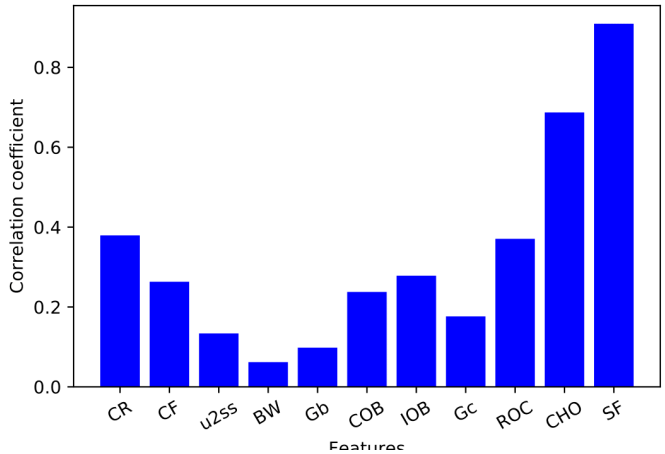

Figure 3.5: The absolute value of the correlation coefficient between the target and the single features.

insulin bolus dose.

3.5 Correlation analysis

In this section, a correlation analysis to assess the relationship between the extracted features and the target was performed. In particular, before starting the creation of the regression models, we checked the entity of the correlation both between the selected features and the optimal bolus and between each feature, to state if multicollinearity exists within the covariates.

3.5.1 Correlation among features

In machine learning, multicollinearity occurs when two or more explanatory variables, which are assumed to be independent of each other, are revealed to be closely related, resulting in redundant features. Hence, in this section, we investigated the entity of multicollinearity within the generated dataset. Figure 3.6, depicts the Pearson correlation coefficients (ρ) computed between each feature. In particular, the features highly correlated each other are CF and Ib

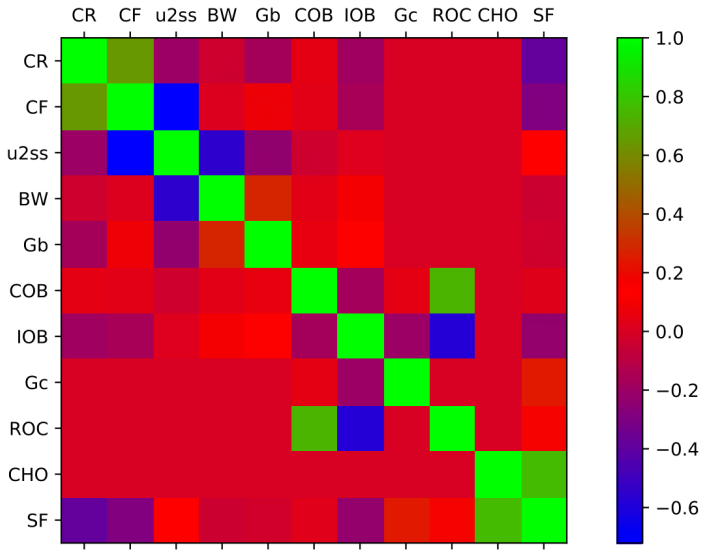

Figure 3.6: Map reporting the correlation between the features.

(ρ = -0.73), CR and CF (ρ = 0.65), ROC and IOB (ρ = -0.58), ROC and COB (ρ = 0.74) BW and Ib (ρ = -0.55). As expected, SF had nonzero correlation with the majority of the variables, especially with CHO (ρ = 0.74), CR (ρ = -0.41) and CF (ρ = -0.28). All the other correlation coefficients are lower compared to the mentioned ones, despite, as expected, all the variables having a nonzero correlation, although low, with another variable. Hence, strong multicollinearity is present within the dataset, leading to a possible loss in reliability when determining the effects of the individual features on the dependent variable within the machine learning model. Indeed, multicollinearity could lead to interpretability issues when analysing the regression coefficients of a linear model. Hence, the effect of multicollinearity should be taken into account during model development, justifying the use of the LASSO methodology in the following Chapter [68].

3.5.2 Correlation between the features and the optimal bolus dosage

In this section, we checked the correlation of the extracted features and the optimal bolus, to assess whether the features were suitable for predicting the target variable. As expected, the results showed that the most correlated feature was SF, with a Pearson correlation coefficient of $\rho = 0.91$, followed by CHO ($\rho = 0.69$). Also CR and ROC resulted correlated with the target, with $\rho = -0.41$ and $\rho = 0.39$ respectively. Figure 3.5 depicts the absolute values of correlation coefficients through a bar plot.

In conclusion, significant information can be extracted from this correlation analysis, i.e. the target is highly correlated with several variables, suggesting that these variables (CHO, SF, CR and ROC) are relevant regressors for the models.

3.6 Comparison between the optimal bolus and the SF bolus dosage

To perform a comparison between the optimal bolus and SF performances, different representative BG curves related to the application of both the optimal insulin bolus and SF are reported. Figure 3.7 shows a representative scenario in which the amount of injected insulin suggested by the SF (1.6610 U) leads to a poorer glycemic control compared to the optimal bolus (3.1157 U), in particular, a hyperglycemic event occurs when injecting SF after the meal, while the application of the optimal bolus allows a tight glycemic control. In Figure 3.8 a hypoglycemic event following the SF dose injection (4.7028 U) is depicted, indeed the mealtime insulin dosage suggested by SF is higher than the optimal one (3.7069 U) leading to glucose levels below the hypoglycemic threshold.

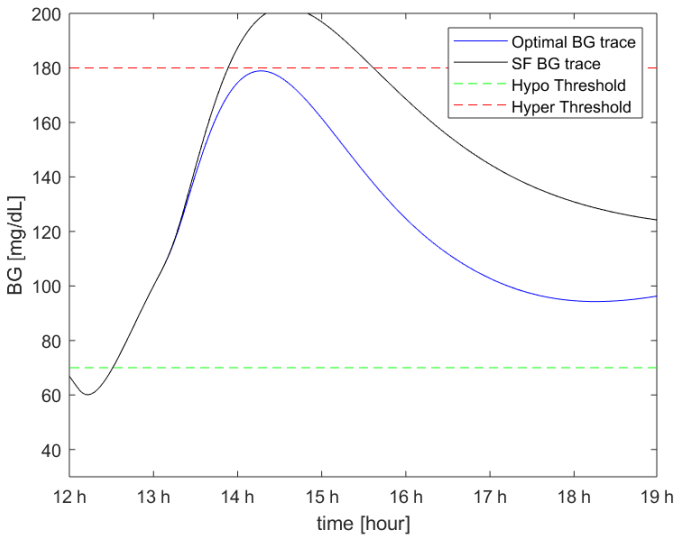

Figure 3.7: Representative BG curve of subject 64, having the following mealtime conditions at 1 pm: BG=100 mg/dl, ROC=1 mg/dl/min, CHO=50 g, Bopt=3.1157 U/min, SF=1.6610 U/min

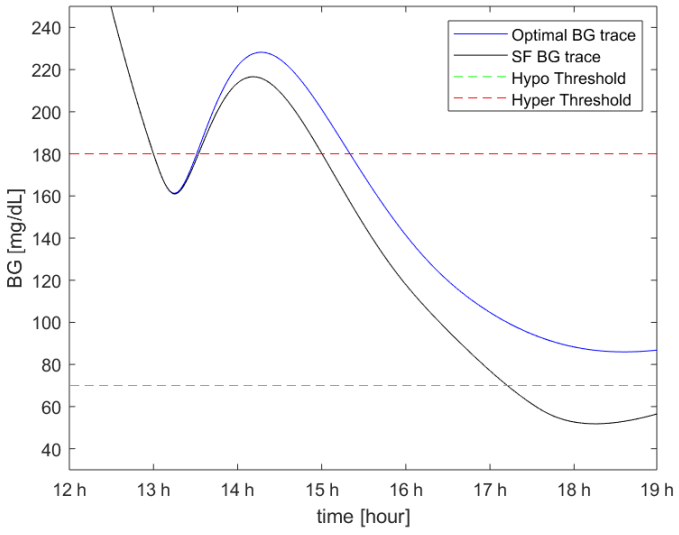

Figure 3.8: Representative BG curve of subject 64, having the following mealtime conditions at 1 pm: BG=180 mg/dl, ROC=-2 mg/dl/min, CHO=120 g, Bopt=3.7069 U/min, SF=4.7028 U/min.

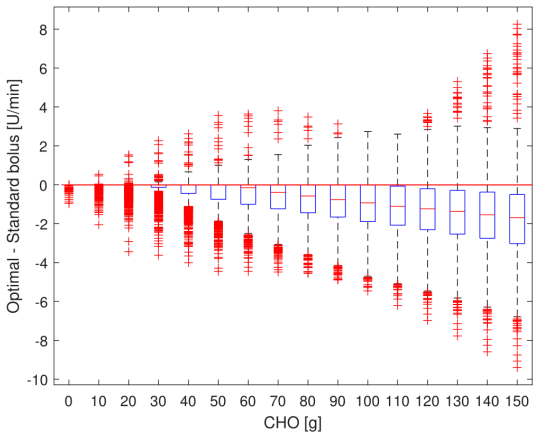

Figure 3.9: Optimal bolus vs Standard bolus distribution depending on CHO values with fixed ROC = -2 mg/dL/min

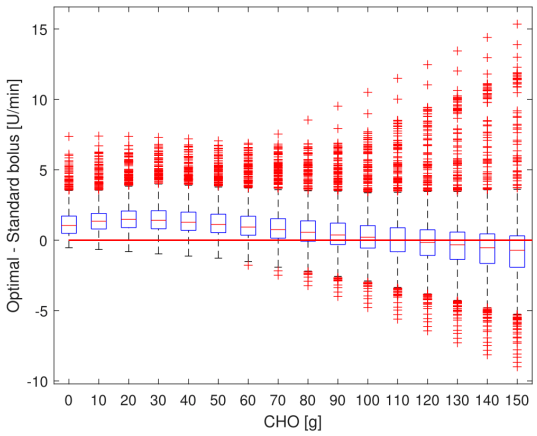

Figure 3.10: Optimal bolus vs Standard bolus distribution depending on CHO values with fixed ROC = 0 mg/dL/min

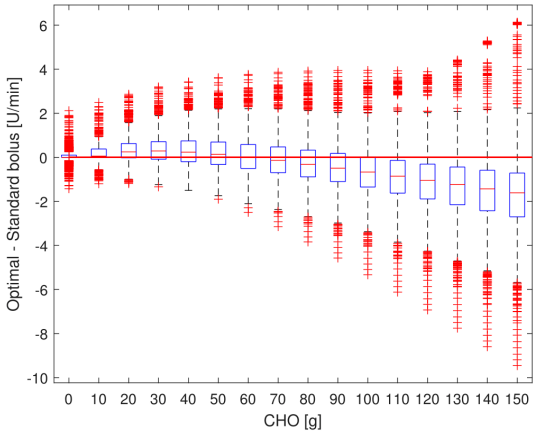

Figure 3.11: Optimal bolus vs Standard bolus distribution depending on CHO values with fixed ROC = 0 mg/dL/min

In Figures 3.9, 3.10, 3.11, three distributions of the difference between optimal and standard insulin bolus based on the CHO amount of the meal are depicted. By analysing the median of the distribution in Figure 3.9, we found out that for higher amounts of CHO, the standard bolus size systematically overestimates the optimal reference. On the other hand, in Figure 3.10, it may be noted that when the ROC is stable, the bolus dosage is underestimated for small quantities of CHO, while when the ROC is rising (Figure 3.11) the insulin dose is moderately underestimated for low CHO values and overestimated for high values.

In conclusion, this analysis highlighted the limitations which characterize the standard bolus dosage, especially for specific mealtime conditions of the subject. An improper bolus amount could lead the patient to hyper- or hypoglycemia, hence the need to improve this calculation. In the following Chapter will be proposed new mealtime insulin bolus models, based on machine learning techniques, to achieve better glycemic control and to overcome the aforementioned limitations.

Chapter 4

Linear and nonlinear supervised learning models for insulin dosing

The real-time availability of information on glucose dynamics provided by CGM systems, along with the possibility of using the BG measurements they provide for insulin dosing, has encouraged the development of new rules to adjust the SF according to data provided by such devices.

The derivation of the literature methods aimed at adjusting the SF discussed in Chapter 2 has mainly been empirical, suggesting that there would be room for improvement should a systematic modelling methodology be adopted. Moreover, in our recent in silico assessment of the state-of-art adjustments we showed that all methods performed similarly for all BG and ROC conditions at mealtime [46], encouraging the research for more effective, and possibly personalized, insulin bolus calculation strategies. Some recent proof-of-concept studies showed that supervised learning techniques can be used to tailor the SF correction to meet patient-specific parameters [44]. However, such methodologies could even be employed to design new models for IB estimation, that is, by abandoning the idea of using the SF as an initial estimate to be adjusted according to ROC, as performed in the previous works, and developing a novel

approach which integrates personalized parameters and CGM dynamics information.

Thus, the study presented in this Chapter aims at developing new approaches to mealtime insulin dose calculation which are based not only on the parameters already included in equation (1.1), but also on the glucose ROC that is provided by CGM and other easily accessible patient-dependent variables, by applying different supervised learning techniques while increasing model complexity.

The work presented in this Chapter was published in the papers of Noaro et al. [48, 49, 50].

4.1 Supervised learning framework

Dataset: The proposed models were developed by employing the simulated dataset generated in the previous Chapter, which is composed of different mealtime conditions related to 100 adult virtual subjects. Starting from the aforementioned database, we excluded all the records corresponding to an amount of CHO equal to zero grams, being this work focused on the investigation of a model to target the mealtime insulin bolus, thus requiring the occurrence of a meal within the scenario. Moreover, we discarded the feature related to the COB, which, compared to the other variables, resulted not easily accessible during a real-life application of the proposed model, since it requires the log of all the meals and snacks within the preprandial 6-hour time window, increasing the burden of the therapy.

In addition, during the first phase of this work, i.e., during the development of linear models, we employed dataset records having prandial BG levels within the normoglycemic range, i.e. from 70 to 180 mg/dL, as described in Chapter 3; while, at a later stage, i.e., during the development of nonlinear ap-

proaches, we decided to extend the dataset, by integrating mealtime scenarios having prandial BG measurements also above this range, that is, mealtime BG ranging from 70 to 250 mg/dL. This choice was made to design a model having a wider validity domain since the consumption of meals in hyperglycemia is a frequent occurrence in T1D and being the nonlinear approaches more suitable to handle a wider domain of the input features.

Thus, for each model, the resulting simulated dataset was divided into training and testing sets. The data on 80 subjects were assigned to the training set, while the remaining data, on 20 subjects, were assigned to the test set. The assignment of each virtual subject either to the training or to the test set was performed randomly. Note, also, that each subject was included either in the training or in the testing set, to provide an unbiased evaluation of model performance within the testing set, composed of distinct subjects. Then, variables composing the dataset were standardized and scaled to unit variance due to the different measurement scales among the features.

Hyperparameters tuning: Both the linear and the nonlinear approaches were developed within the same machine learning framework. In particular, when the model structure included any hyperparameter to be tuned, we resorted to a 5-fold cross-validation in the training set. This implies that the original training set is partitioned into five equally sized subsets. Then 5 iterations are performed, where one of 5 subsets is retained as the validation set, while the others are used to train the model with the selected hyper-parameters configuration. This procedure is repeated for each fold, and the returned scores are averaged. To tune the hyper-parameters, the 5-fold cross-validation is applied over a fixed parameter grid. In the end, the set of tuning parameters for which the 5-fold cross-validation reports the lowest error is chosen to train the final model.

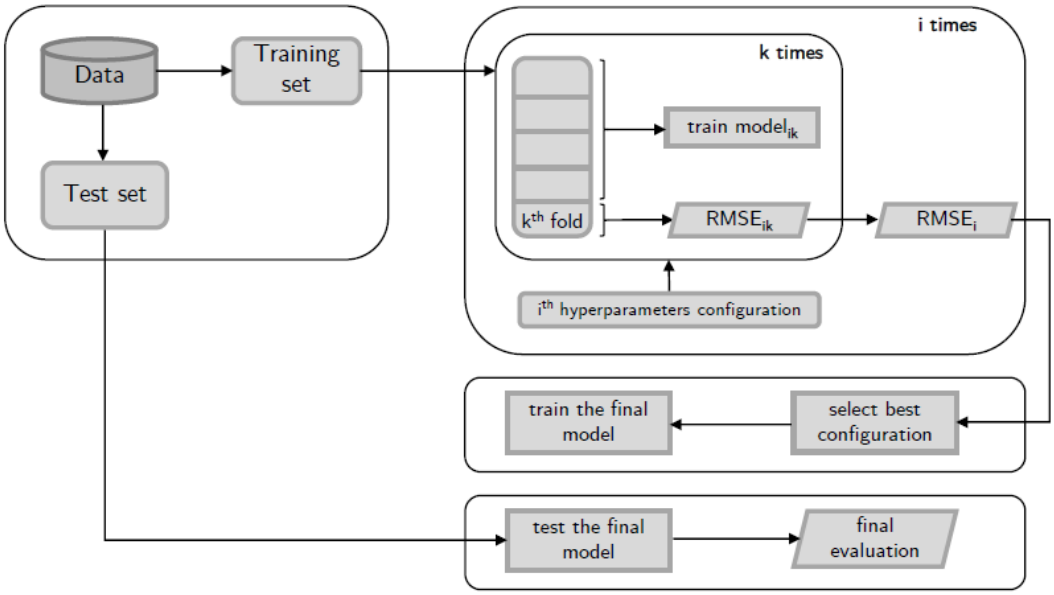

Figure 4.1: Cross-validation framework employed for the hyperparameters tuning of the proposed models.

4.2 Linear regression approaches

During the first phase of this study, we selected multiple linear regression (MLR) and least absolute shrinkage and selector operator (LASSO) [69] among the possible approaches to target the optimal mealtime insulin bolus (IB_{OPT}) being these methodologies simple and able to provide an adequate and interpretable description of how the inputs affect the output. Indeed, each MLR coefficient represents the slope of the linear relationship between the output and that portion of input which is independent of all the others. Moreover, model interpretability represents a desirable feature for clinicians, which could encourage them to use it in clinical practice.

4.2.1 Multiple linear regression (MLR)

The MLR model has the following form:

$$\hat{\boldsymbol{y}} = \hat{\alpha}_0 + \sum_{j=1}^{p} \boldsymbol{x}_j \cdot \hat{\alpha}_j \tag{4.1}$$

where $\boldsymbol{y}$ is the target variable, i.e. IB_{OPT}, $\boldsymbol{x_j}$ is the j-th feature, α_j the coefficient related to the j-th feature, α_0 is the model intercept and p represents the number of features. Parameters $\hat{\alpha}_j$ are estimated through the least squares estimation method [69], which chooses a vector $\hat{\boldsymbol{\alpha}}$ of coefficients that minimizes the residual sum of squares (RSS):

$$\hat{\boldsymbol{\alpha}} = \underset{\boldsymbol{\alpha}}{\operatorname{argmin}}\, RSS(\boldsymbol{\alpha}) \tag{4.2}$$

where

$$RSS(\boldsymbol{\alpha}) = \sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \tag{4.3}$$

and y_i is the i-th observation of IB_{OPT}, $\hat{y}_i$ the corresponding model prediction.

the resulting MLR equation, identified on the training set, is reported in eq. (4.5). As expected CHO, SF, and ROC contribute positively to the final insulin amount. Moreover, their coefficients are generally bigger than the others (absolute value), thus highlighting the importance of these features for IB computation. CR makes a negative contribution since the lower the CR the higher the amount of insulin required to compensate for a specific CHO intake. Similar reasoning, but in terms of insulin sensitivity, can be applied to CF, which has a negative sign. On the other hand, IOB and BW coefficients present positive and negative signs, respectively, which are the opposite of those expected from a physiological interpretation of these variables. This result is probably due to the presence of multicollinearity among features (as reported in Section 3.5.2).

4.2.2 Least absolute shrinkage and selector operator (LASSO)

To deal with multicollinearity, we employed a shrinkage method. In particular, we resorted to the LASSO regression technique, which is well known to be robust to multicollinearity [68]. The LASSO coefficients are estimated by minimizing eq. (4.3)) with the addition of the absolute value of the magnitude of the coefficient as a penalty term:

$$\hat{\boldsymbol{\alpha}} = \underset{\boldsymbol{\alpha}}{\operatorname{argmin}}\{RSS(\boldsymbol{\alpha}) + \lambda \sum_{j=1}^{p} |\alpha_j|\} \tag{4.4}$$

where $\lambda \geq 0$ is a parameter controlling the amount of shrinkage set, through an exhaustive grid search, with cross-validation in the training set. In this work, we trained three LASSO models on three different feature sets:

- LASSO: trained on the feature set described in Chapter 3 defined as $\{x_j : j = 1, ...p\}$.

- LASSO_Q: trained on an expanded feature set which includes variables reported in Chapter 3 plus their quadratic values, defined as $\{x_j, x_j^2 : j = 1, ...p\}$.

- LASSO_{QI}: trained on a extended feature set which also includes terms of between-features interaction, defined as $\{x_j, x_j^2, x_{ij} : j = 1, ...p, i = 1, ...p, i \neq j\}$.

Note that, due to the intrinsic nonlinearity of the glucose-insulin system, we also added polynomial transformations of the input variables as features, thus capturing nonlinear relationships between variables while still maintaining model interpretability.

One of the key features of the LASSO model is that it performs both automatic variable selection simultaneously, setting the coefficients associated with

unnecessary features to zero, and, also regularization. In practice, this means that the LASSO model slightly increases the bias to reduce the variance of the predicted values: this leads to an overall improvement in the accuracy of predictions [69]. In this application here, this "natural" feature selection capability made it possible to considerably reduce the number of features, and to avoid overfitting, especially when quadratic (LASSO_Q) and quadratic plus interaction (LASSO_{QI}) terms were included in the dataset.

LASSO

The LASSO equation identified is reported in eq. (4.6). Note that variables CF, Ib and IOB were discarded during the LASSO training procedure, by adopting the automatic selection feature offered by this methodology. This result was expected, since correlation analysis had revealed a high correlation between these features and CR, BW, and ROC, respectively. Note also that the CR, BW and ROC coefficients had changed in (absolute) magnitude when compared to those of the MLR model. The CR coefficient, in particular, has increased, whereas the BW and ROC coefficients have decreased, as well as the Gc coefficient.

LASSO_Q

The final LASSO_Q equation is reported in eq. (4.7). Note that, by adding the quadratic terms, only the most relevant first-order features (CR, ROC, CHO, SF) were selected in the training procedure, while BW, Gt and Gc appear within the model only with a quadratic contribution.

LASSO$_{QI}$

The LASSO$_{QI}$ equation identified on the training set is reported in eq. (4.8). More specifically, augmenting the inputs with both quadratic and interaction terms leads to the elimination of all the first-order terms, thus lending more importance to the interaction and quadratic terms. In particular, the highest coefficient is related to (Gt,SF) interaction, followed by (Gt,ROC), while the other coefficients are very close to zero.

Remark: Note that the three LASSO models were trained on three different feature sets. The value of the hyperparameter λ was chosen, for all the models, by using cross-validation, as highlighted in Section 4.1, and particularly by searching, exhaustively, among 200 equally-spaced values ranging between 0.001 and 10 and by selecting the value that maximized the R^2 in a 5-fold cross-validation ($\lambda = 0.05$).

$$IB_{MLR} = 4.603 - 0.137\,CR - 0.191\,CF - 0.181\,I_b - 0.380\,BW + 0.464\,G_t + 0.039\,IOB - 0.065\,G_c + \\ +0.858\,ROC + 0.273\,CHO + 2.686\,SF \tag{4.5}$$

$$IB_{LASSO} = 4.603 - 0.214\,CR - 0.238\,BW + 0.410\,G_t - 0.032\,G_c + 0.806\,ROC + 0.257\,CHO + 2.661\,SF \tag{4.6}$$

$$IB_{LASSO_Q} = 4.603 - 0.198\,CR + 0.789\,ROC + 0.234\,CHO + 2.671\,SF - 0.224\,BW^2 + 0.403\,{G_t}^2 - 0.020\,{G_c}^2 \tag{4.7}$$

$$IB_{LASSO_{QI}} = 4.603 - 0.150\,CR \cdot G_t - 0.064\,CR \cdot G_c + 0.007\,CF \cdot SF + 0.043\,I_b \cdot ROC + 0.064\,I_b \cdot CHO + \\ +0.019\,I_b \cdot SF - 0.199\,BW^2 + 0.064\,BW \cdot ROC + 0.236\,{G_t}^2 + 0.632\,G_t \cdot ROC + \\ +0.167\,G_t \cdot CHO + 2.66 G_t \cdot SF + 0.014\,IOB \cdot SF + 0.079 ROC \cdot SF \tag{4.8}$$

4.2.3 Models evaluation and results

Error in estimating the optimal amount of insulin

The aim of this first evaluation, carried out on the simulated test set, was to assess and quantify whether the models developed would be able to estimate IB_{OPT} more accurately than SF, BU, SC, and ZI. This was carried out by computing both the root mean square error (RMSE) and the coefficient of determination (R^2), between the optimal and the estimated IB for each model. Table 4.1 reports the results obtained in terms of RMSE and R^2. Models MLR, LASSO, $LASSO_Q$ and $LASSO_{QI}$ estimate the optimal insulin bolus more accurately when compared with the other methods. Specifically, RMSE is 1.45 U for SF, 1.36-1.44 U for the literature models, and 0.84-0.87 U for the new models, with the best result achieved by $LASSO_{QI}$ (RMSE = 0.84 U).

The R^2 metric also improved with the new models (R^2 = 0.91-0.92), achieving the highest value with $LASSO_{QI}$ (R^2 = 0.92) when compared with SF (R^2 = 0.82) and with the methods described in the literature (R^2 = 0.84-0.85). Note too, that BU, SC and ZI also slightly improved their performances when compared with SF, but both lower RMSE and higher R^2 values were obtained with the proposed new models.

Table 4.1: Comparison of metrics for prediction accuracy and goodness of fit evaluation. Values related to SF, state-of-art methods and the models proposed are reported.

Metric	SF	BU	SC	ZI	MLR	LASSO	$LASSO_Q$	$LASSO_{QI}$
$RMSE$ [U]	1.45	1.44	1.36	1.44	0.87	0.87	0.86	0.84
R^2	0.82	0.84	0.85	0.85	0.91	0.91	0.91	0.92

Table 4.2: Comparison of metrics assessing glycemic control for SF, state-of-the-art methodologies and the models proposed. Metrics related to IB_{OPT} are reported as reference values. Median and interquartile ranges are reported for TIR, TBR, TAR, and BGRI.

Metric	IB_{OPT}	SF	MLR	LASSO	$LASSO_Q$	$LASSO_{QI}$	BU	SC	ZI
TIR %	68.14 (59.00-82.27)	60.11 (38.23-78.12)	64.54 (47.92-78.94)	64.55 (47.92-79.10)	64.55 (48.20-79.22)	64.82 (49.58-79.78)	62.04 (37.39-78.94)	62.05 (37.95-78.50)	60.66 (36.56-78.67)
TBR %	0.00 (0.00-0.00)	0.00 (0.00-28.53)	0.00* (0.00-14.68)	0.00* (0.00-13.99)	0.00* (0.00-13.57)	0.00* (0.00-10.25)	0.00 (0.00-28.81)	0.00 (0.00-28.25)	0.00 (0.00-28.81)
TAR %	29.92 (12.19-39.34)	29.09 (13.30-37.95)	29.92 (12.19-39.06)	29.92 (11.91-39.06)	29.92 (11.90-39.34)	30.19 (11.63-39.34)	29.09 (13.60-37.67)	29.09 (13.30-37.81)	29.36 (14.13-37.95)
BGRI	8.23 (3.97-14.03)	9.93 (4.85-17.46)	9.10* (4.68-15.50)	9.09* (4.67-15.45)	9.08* (4.68-15.43)	8.97* (4.59-15.31)	9.53 (4.67-17.34)	9.72 (4.74-17.28)	9.68 (4.69-17.76)

[1]*Statistically significant compared to SF with p-value < 0.0071 (Bonferroni-corrected threshold)

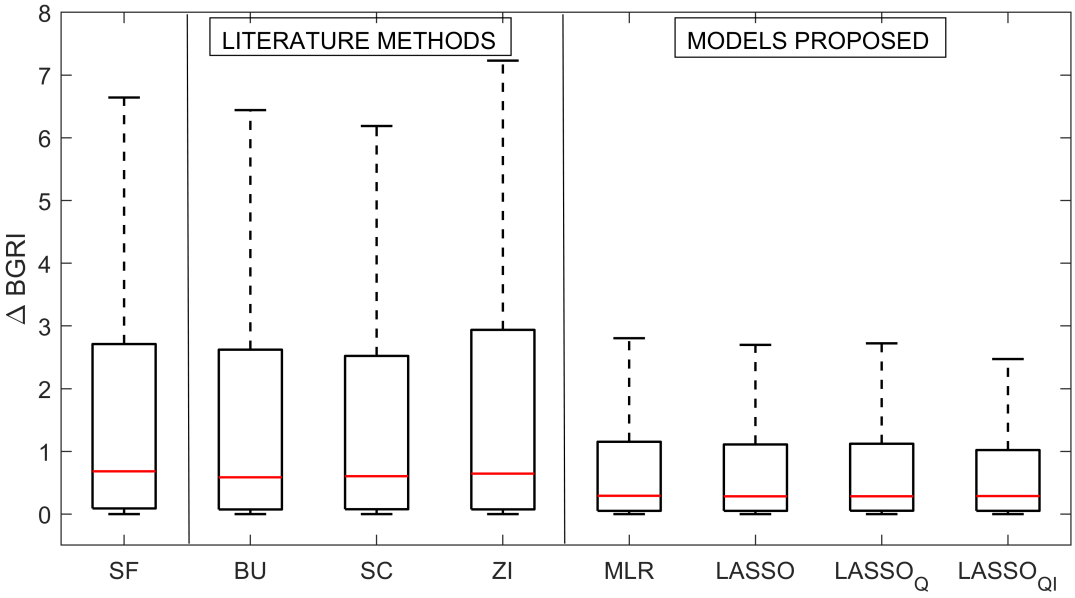

Figure 4.2: Distribution of the difference between BGRI of SF, MLR, LASSO, LASSO$_Q$, LASSO$_{QI}$, BU, SC, ZI methods versus IB$_{OPT}$.

Assessment of glycemic control

The MLR and LASSO models were also compared against SF, BU, SC and ZI in terms of glycemic outcome. In particular, the percentage of TIR, TBR, TAR, and the BGRI, explained in Chapter 2 were computed. The results obtained are reported in Table 4.2. When considering the BGRI, SF revealed the highest median risk (9.93) among all the other methods, followed by BU, SC and ZI which showed a median BGRI of 9.53, 9.72 and 9.68 respectively. The lowest median BGRI values were obtained by LASSO$_Q$ (9.08) and LASSO$_{QI}$ (8.97), both of which were close to the BGRI value that was obtained using MIB$_{OPT}$ (8.23). Fig. 4.2 shows the distribution of the difference in BGRI (ΔBGRI) between each of the IB calculation methods and IB$_{OPT}$. Since the optimal insulin bolus minimizes the BGRI function, the ΔBGRI distributions are in the positive half-plane. Note that the BGRI distribution of our models is closer to that of IB$_{OPT}$ (lower ΔBGRI values) compared to the other methodologies.

In terms of TIR, all the proposed methods outperformed both the SF and

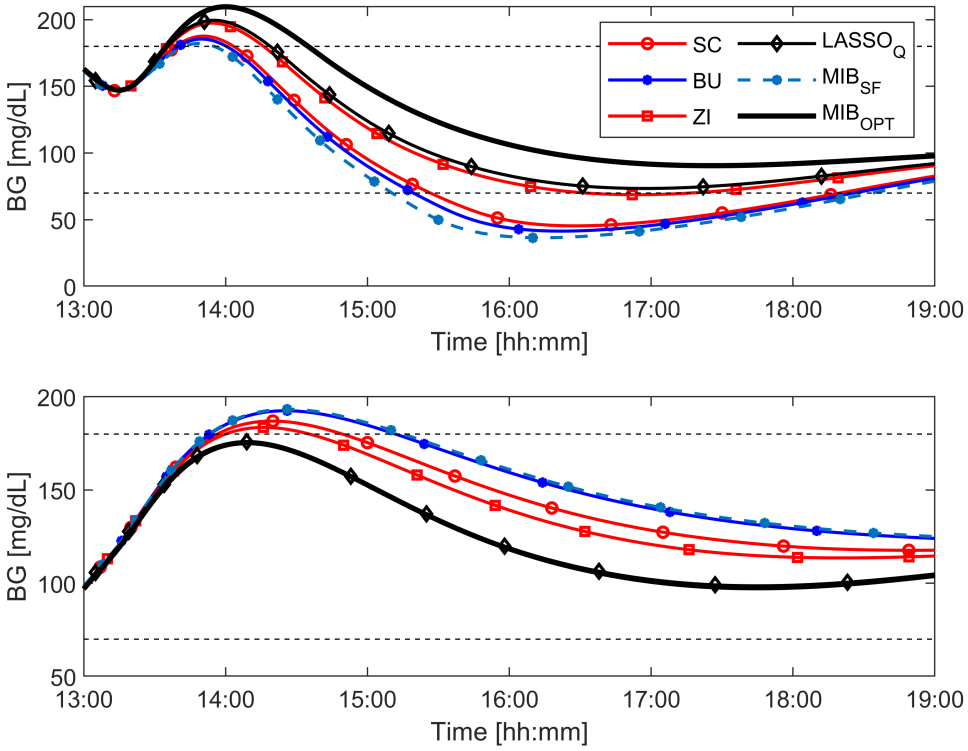

Figure 4.3: Representative examples of BG curves during postprandial time window for different methods of insulin bolus computation and different mealtime conditions. For better visualization, only LASSO$_Q$ among the models proposed is reported. In the upper panel, mealtime ROC is negative (-1.5 mg/dL/min), starting BG=160 mg/dL and meal CHO is 60 g. The calculated IB doses are IB$_{OPT}$=1.89 U,IB$_{SF}$=3.62 U, IB$_{LASSO_Q}$=1.94 U, IB$_{BU}$=3.26 U , IB$_{SC}$=3.02 U, IB$_{ZI}$=2.12 U. In the lower panel, mealtime ROC is positive (1.5 mg/dL/min), starting BG=100 mg/dL and meal CHO is 30 g. The MIB doses are IB$_{OPT}$=2.80 U, IB$_{SF}$=0.72 U, IB$_{LASSO_Q}$=2.74 U, IB$_{BU}$=0.79 U, IB$_{SC}$=1.34 U, IB$_{ZI}$=1.71 U. Dashed lines indicate the euglycemic range.

the literature methods, while, as regards the TBR metric, the median values proved not to be informative as they were equal to 0 for all methods. On the other hand, when considering the 75^{th} percentile of TBR, it could be stated that the magnitude of hypoglycemic events is considerably reduced for the proposed models compared to SF, BU, SC and ZI. The 75^{th} percentile of TBR in particular, decreases from about 28% with SF, BU, SC and ZI to a value between 10.25% and 14.68% obtained with the new models. The best outcome was achieved by LASSO_{QI}, with a 75^{th} percentile of TBR equal to 10.25%. In addition, the improvement in terms of BGRI and TBR given by MLR, LASSO, LASSO_{Q} and LASSO_{QI} is statistically significant (p-value < 0.0071) when compared to SF. Regarding hyperglycemia, the median TAR values slightly increased for all new models when compared with the literature methods. This result was expected since the BGRI, which is the cost function minimized to compute IB_{OPT}, assigns a higher risk to hypoglycemia than it does to hyperglycemia, thus resulting in higher TAR values for IB_{OPT} (and, consequently, also for the models proposed, which target IB_{OPT}) when compared to the methodologies proposed in the literature. However, the TAR increase is moderate and is not statistically significant when compared with SF. Moreover, it does not negatively affect overall glycemic control in terms of BGRI. Figure 4.3 shows two representative postprandial BG curves, after the administration of IB, computed through one of the candidate models, i.e., LASSO_{Q}, IB_{OPT} and the literature methods. Note that, only LASSO_{Q} has been considered as a representative model, for reasons of better visualization. As shown in the upper panel, SF, BU and SC all induce hypoglycemia, while LASSO_{Q} and ZI allow a proper glycemic control, approaching that of IB_{OPT}, despite the initial hyperglycemia that was mainly due to the meal, and to high pre-prandial BG. The lower panel shows the occurrence of hyperglycemic events after the application of SF, BU, ZI and SC methodologies, while LASSO_{Q} led to optimal glycemic control.

4.3 Nonlinear approaches

To extend the work presented within the previous section, we tested three candidate nonlinear models targeting the IB_{OPT}, trying to capture the nonlinear connections between the features, being the nature of the glucose-insulin system nonlinear. Specifically, we tested models based on an ensemble of multiple regressors, which is a machine learning approach aimed at combining multiple models within the learning procedure. Such a model could be implemented leveraging diverse methodologies, such as bootstrap aggregating, gradient boosting, or in the form of a voting procedure. In this work, we first tested two different ensemble learning models for regression, which take advantage of the bootstrap aggregating and the boosting strategies, that are random forest (RF) and gradient boosted tree (GBT) algorithms respectively [69]. In the second stage, to take advantage of the distinct predictive abilities which characterize both a linear and a nonlinear approach, we designed a dynamic voting ensemble algorithm, by combining a linear, i.e. LASSO, and a nonlinear, i.e. RF, model. Lastly, the models were evaluated in-silico, as performed for the linear regression models.

4.3.1 Random forest (RF) and gradient boosted tree (GBT) models

In this section, we discuss two tree-based models used to predict the optimal insulin dose at mealtime. Particularly, the two methods, i.e. RF and GBT, take advantage of two common strategies in ensemble learning, which are bootstrap aggregating and boosting algorithms. In the following, the two methods will be outlined.

Random forest model

The RF model leverages the general technique of bootstrap aggregating, i.e. bagging, which is an ensemble-based algorithm. In particular, given a training dataset D of cardinality N, the bagging algorithm trains T independent classifiers, each trained by sampling, with replacement, N instances (or some percentage of N) from D. Bagging has been proven to outperform especially for high variance and low bias procedures, such as trees. Indeed, the RF model is a modification of bagging, where an ensemble of decision trees is generated by training the models following the bootstrap aggregated strategy. In addition to choosing instances, RF uses a modified tree learning algorithm that incorporates a randomized feature selection. Hence, the final prediction is given by the aggregation of the trained weak learners by averaging their outputs [70, 71]. In this work, we tuned through the cross-validation procedure described in section 4.1 the following hyperparameters of the RF model: the number of trees, the number of features considered at each splitting node and the number of levels in each decision tree.

Gradient boosted tree model

On the other hand, the GBT learning algorithm leverages the gradient boosting technique. The boosting technique, just like bagging, is an iterative approach for generating a strong classifier, capable of building a strong model by exploiting an ensemble of weak learners. However, it differs from bagging in one crucial way, since it relies on an ensemble of individual learners, which are not trained independently, as in the bagging technique, though sequentially. Thus, in boosting, the training dataset for each subsequent classifier increasingly focuses on instances misclassified by previously generated classifiers. In gradient boosting, which is a modification of boosting, the main idea

is to construct the new individual learner to be maximally correlated with the negative gradient of the loss function, associated with the whole ensemble. Hence, this method fit the new learner to the residual errors, which can be interpreted as negative gradients, made by the previous learner. Also in this case, when the weak learner is selected as a decision tree model, the resulting algorithm is called gradient boosted tree [72]. In this work, we tuned through the cross-validation procedure described in section 4.1 the following hyperparameters of the GBT model: the maximum depth of a tree, the minimum sum of weights of all observations required in a child, the minimum loss reduction required to make a split, and the L2 regularization term on weights.

4.3.2 Models evaluation and results

The evaluation of the models was performed on the testing set. First, we assessed the accuracy in targeting the IB_{OPT} by computing the RMSE between IB_{OPT} and the IB estimated by the methods proposed.

Then, the two models were compared in terms of postprandial glycemic control, using the UVA/Padova T1D Simulator. In particular, the IB estimates were tested with regard to the glycemic outcome, by computing the TIR, TAR and TBR, and BGRI metrics. The two proposed models proposed were compared to the best performing linear model in terms of overall glycemic risk metrics, i.e. BGRI, presented in the previous Section, that is the $LASSO_{QI}$ model, and the statistical significance of the differences of the metrics distributions was evaluated by applying Friedman test, and correcting the significance level for multiple comparisons.

The $LASSO_{QI}$, RF and GBT $RMSE$ in estimating IB_{OPT} are 0.92 U, 1.00 U, 0.91 U, respectively. These first results may suggest that the difference between $LASSO_{QI}$, RF and GBT models seems almost negligible. However, when com-

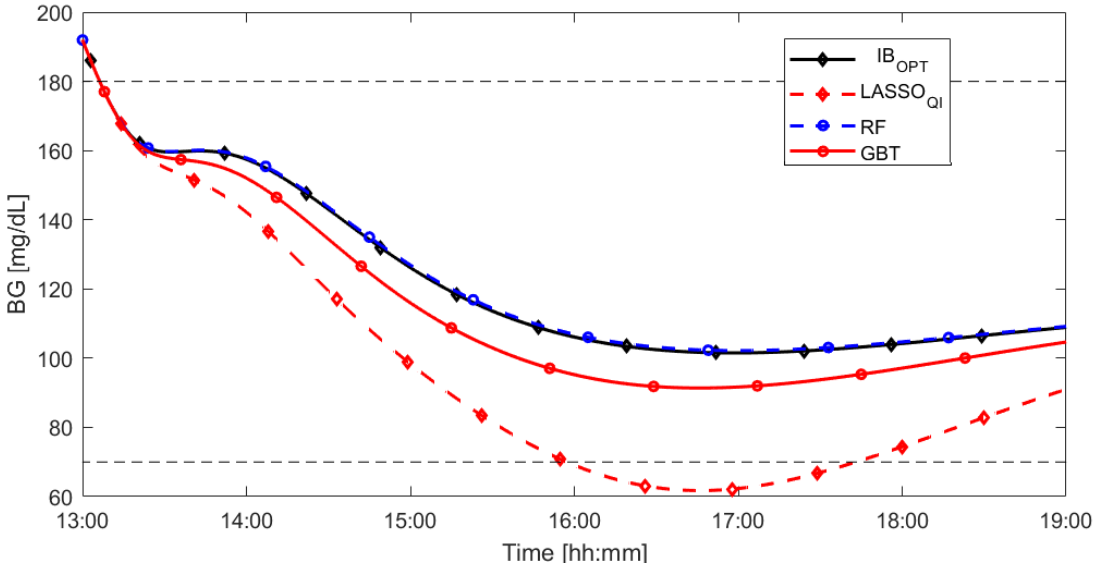

Figure 4.4: Representative example of BG curves during postprandial time window for different methods of insulin bolus computation. The prandial ROC is negative (-2 mg/dL/min), starting BG=190 mg/dL and meal CHO is 20 g. The calculated IB doses are IB_{OPT}=0.72 U, $IB_{LASSO_{QI}}$= 2 U, IB_{RF}=0.69 U , IB_{GBT}=0.98 U.

paring the two nonlinear models and analysing the absolute error of each estimate, GBT proved to be better than RF in targeting the IB_{OPT} for meals with CHO content higher than 40 g. For those meals, the $RMSE$ achieved by RF is equal to 1.11 U, while that of GBT is 0.98 U. On the other hand, the RF model shows better accuracy in the estimation of IB_{OPT} for meals with CHO content lower than 40 g ($RMSE$ achieved by RF and GBT are 0.49 U and 0.53 U, respectively).

From the glycemic control point of view, results reported in Table 4.3 show that, when both GBT and RF are used for insulin dosing, median BGRI improves. Specifically, the BGRI values obtained by $LASSO_{QI}$, GBT and RF are 10.38, 9.86, and 10.03, respectively, and only the reduction achieved with GBT is statistically significant. Moreover, it is possible to appreciate the reduction of the 75^{th} percentile of TBR from 13.02 % of $LASSO_{QI}$ to 0.28 and 3.88 % for GBT and RF, respectively. Finally, it is worth underlining that the TBR distribution of both the nonlinear models is significantly different from the one of

Table 4.3: Quantitative assessment of glycemic control for $LASSO_{QI}$, RF and GBT: the median and interquartile range is reported for TIR, TBR, TAR and BGRI. Values in the IB_{OPT} column are those obtained with the optimal insulin bolus, which should be assumed as a reference.

Metric	IB_{OPT}	$LASSO_{QI}$	RF	GBT
TIR %	65.10 (55.67-76.73)	60.38 (44.32-74.50)	61.77 (47.92-76.45)	62.05 (49.30-75.62)
TBR %	0.00 (0.00-0.00)	0.00 (0.00-13.02)	0.00* (0.00-3.88)	0.00* (0.00-0.28)
TAR %	33.79 (21.05-42.66)	33.51 (20.22-42.94)	33.52 (19.94-43.49)	34.07 (20.50-43.49)
BGRI	9.23 (4.85-15.15)	10.38 (5.74-16.85)	10.03 (5.24-16.73)	9.86* (5.21-16.20)

[2]*Statistically significant compared to $LASSO_{QI}$

$LASSO_{QI}$. Fig. 4.4 shows four postprandial BG curves obtained as a response to administering IB_{OPT}, $LASSO_{QI}$, RF and GBT insulin boluses in correspondence to the same meal of the same virtual subject. Note that, the IB related to $LASSO_{QI}$ overestimates the optimal dose, leading to a hypoglycemic event. On the other hand, both the nonlinear methods can better estimate IB_{OPT}, ensuring a safer glycemic control. In particular, the RF model estimates more accurately IB_{OPT} than GBT (IB_{OPT} = 0.72 U, IB_{RF} = 0.69 U , and IB_{GBT} = 0.98 U.). This result is not surprising considering that the CHO amount composing the meal is 20 g. Indeed, we have previously highlighted that the RF model showed better accuracy in the estimation of IB_{OPT} for meals with CHO content lower than 40 g, compared to GBT.

4.3.3 Ensemble model based on dynamic voting (DV)

In the previous Sections, both linear and nonlinear models were proposed to target the optimal insulin dosage. However, a further margin of improvement may come from the possibility of combining these two techniques in a single one, thus developing an ensemble model. Indeed, the different outputs can be merged by assigning a higher weight to the model (base learner) which outper-

forms the other within a specific "local region", i.e. within the subspace where the regression error associated with that base learner is the lowest, thus taking advantage that each model produces better results inside certain subareas of the application domain. Following this rationale, in this section, we proposed an ensemble model based on a dynamic voting procedure (DV), thus combining both a linear and a nonlinear model, i.e. LASSO and RF respectively. Therefore, the final prediction is derived by weighting the two outputs according to the local performance, to develop a more effective model for the estimation of the optimal IB dose.

Models composing the ensemble

We selected as base regressors, i.e., as models that composed the ensemble, the LASSO model and the RF model, previously explained. We selected both a linear and a nonlinear model to take advantage of the different predictive abilities of the two models. Indeed, the LASSO model describes the linear relationship between the features and the output, by allowing also for both variable selection and regularization, through the addition of a penalty term to the loss function. On the other hand, the RF model, which is a nonlinear learning method based on an ensemble of decision trees, is focused on capturing nonlinear connections between the features.

Ensemble method implementation

The ensemble method is mainly based on the DV algorithm described in [73], where multiple regressors are combined using local performance estimates to generate the final prediction. In particular, the algorithm is divided into two different phases, described below.

Learning phase: First, the hyperparameters of LASSO and RF were optimized in a cross-validation setup by performing an exhaustive search over a

fixed grid of parameter values. Then, we computed the estimation error for each data record (EE_R) and each base learner in the training set, in another cross-validation setup, as:

$$EE_R = |\widehat{IB}_{lr} - IB_{OPT_r}| \tag{4.9}$$

where $\widehat{IB}_{lr}$ is the output obtained from the l-th learner (fitted on the k-1 training folds of the cross-validation setup) applied to the r-the record of the k-th validation fold, and IB_{OPT_r} is the IB_{OPT} related to the r-the record. Finally, the two base learners, having the optimal hyperparameters previously determined, were trained using the whole training set.

Application phase: The model trained in the previous phase is applied as follows. For each input record of the test set, similar data through the training records were searched using the k-nearest neighbours (k-NN) method, to define its local region [74]. The similarity measure is based on the weighted Euclidean distance, in which weights are derived from the application of the RReliefF (RRF) algorithm to the dataset [75, 76]. This procedure was performed to differently weigh the variables, being some features more relevant than others for the definition of the local subspace, i.e., more related to the target variable. Hence, after having defined the local subarea, for each base learner we predicted the $\widehat{EE}_R$ of the test record by averaging the EE_R of the nearest neighbours. Once the $\widehat{EE}_R$ of both LASSO and RF was derived, the final output was computed by weighting more the prediction of the best performing model than the other, according to $\widehat{EE}_R$. In particular, the prediction related to the model with the best outcome within the local region was weighted three times more than the one having the lowest performance.

4.3.4 Model evaluation and results

We evaluated the efficacy of the ensemble on the test set, both in terms of regression error, i.e., RMSE, and glycemic outcome within the simulated environment, where the estimated IB was applied as mealtime insulin amount and the TIR, TBR, TAR, and BGRI extracted. In addition, the assessment was performed for the single base learners (LASSO and RF), to verify the benefit resulting from the usage of the ensemble compared to the single models. Finally, as a reference, we included the IB_{OPT} within the in silico evaluation. The ensemble method led to an improvement compared to the single models in terms of regression error. In particular, the RMSE resulted in 0.93 U, and 1 U for the base regressors LASSO and RF, respectively, while the DV method produced an RMSE of 0.88 U. Moreover, we found out that the most important subdomain that allows differentiating the performance of the base learners is given by SF. Indeed, within the local subdomain defined from medium SF values (from 8 to 15 U), the LASSO model produced better results compared to RF, while the latter model outperformed LASSO at the extremities of the SF domain. The DV approach allowed the combination of the positive contributions of LASSO and RF, hence improving the outcome.

Table 4.4 reports the median and interquartile ranges of the distributions of TIR, TAR, TBR and BGRI for each considered method. When comparing LASSO with RF, the former resulted in a lower median TAR, while, conversely, the latter outperformed LASSO in terms of TBR, reaching a 75th of 3.88% against 14.66%. Focusing on the results obtained using DV, allowed us to maintain a high postprandial TIR of 62.55% while reaching a trade-off in terms of TAR and TBR. This result was achieved thanks to the ability of the ensemble approach to select as a major contributor to the outcome the base learner having better performance within specific local domains.

Table 4.4: Quantitative assessment of glycemic control for LASSO_{QI}, RF and GBT: the median and interquartile range is reported for TIR, TBR, TAR and BGRI. Values in the IB_{OPT} column are those obtained with the optimal insulin bolus, which should be assumed as a reference.

Metric	IB_{OPT}	LASSO	RF	DV
TIR %	65.10 (55.67-76.73)	60.03 (45.42-73.40)	61.77 (47.92-76.45)	62.55 (49.60-76.00)
TBR %	0.00 (0.00-0.00)	0.00 (0.00-14.66)	0.00 (0.00-3.88)	0.00 (0.00-1.03)
TAR %	33.79 (21.05-42.66)	34.07 (21.02-42.65)	33.52 (19.94-43.49)	33.45 (20.50-42.90)
BGRI	9.23 (4.85-15.15)	10.52 (5.71-16.95)	10.03 (5.24-16.73)	9.76 (5.12-16.00)

4.4 Choice of the final model for the design of an insulin bolus calculator

4.4.1 Motivation

In the following, we will discuss the selection of the final model for the design of an automated mealtime insulin bolus calculator. In this Chapter, we proposed multiple machine-learning-based models to improve the performances in terms of postprandial glucose control, while increasing model complexity. We obtained positive results in silico, and each model, starting from the linear to the nonlinear ones, improved the overall glycemic control compared to SF. However, the increase in model complexity was not corresponding to a remarkable improvement in terms of performance, probably indicating the limitations of a population model for such a task. Beyond that, to select the final model to be potentially integrated within an automated bolus calculator, we considered another relevant aspect that will be explained in the following. Indeed, we also focused on the clinical acceptability of the proposed model, which could potentially substitute the SF in clinical practice. Hence, the main feature for such a model should be interpretability, that is the understanding

of how the input affects the output variable, i.e. the suggested mealtime insulin dose. Indeed, the more interpretable the model is, the easiest would be to trust this model within a clinical setting. For this reason, we selected as the final model a linear one, as linearity leads to interpretable models, being the output of a linear equation representing the relationship between the features and the outcome. This choice was also supported by the slight improvement in terms of glycemic control obtained by the nonlinear and ensemble models compared to the linear ones. Thus, among the candidate linear models, we selected LASSO_Q as the final model, since it allowed the greatest reduction of BGRI and TBR, while not increasing the median TAR when compared with IB_{OPT}, as in LASSO_{QI}. Moreover, LASSO_Q was chosen over LASSO_{QI} because of the lower number of regressors. Indeed, a high number of features could lead to a model having a higher variance, where a small change in the feature values is reflected in a significant change in the output variable. This aspect was considered as the models were developed within a noise-free scenario, hence a potential real-life application could lead to features having different distributions, being affected by errors. Thus, a model having lower variance is key to preventing an incorrect estimation due to error sources.

Therefore, as all the models were tested within a noise-free scenario, thus providing only a preliminary evaluation of their effectiveness, we employed a simulation tool, i.e. ReplayBG, to evaluate the selected model retrospectively on data collected from real patients affected by T1D.

4.4.2 In-silico clinical trial using ReplayBG

In this Section, we will discuss the assessment of the impact of LASSO_Q on already acquired glucose traces thanks to the ReplayBG simulation tool [77]. Indeed, to conduct this evaluation, we decided to resort to a model-based strat-

egy. The approach which characterizes the ReplayBG simulation tool consists of two main steps: first, a state-of-the-art composite physiological model of glucose-insulin dynamics is identified on each selected glucose trace which follows the meal event; second, the model is used to simulate glucose concentration during the meal by replacing the real injected insulin bolus with the insulin dose provided by LASSO_Q. Hence, in the following, the extraction of the glycemic intervals related to the meal and the application of ReplayBG to the selected glucose data will be outlined.

Extraction of the glycemic intervals: Data collected during a randomised crossover trial in patients with T1D [78] were used for the retrospective analysis to assess the effectiveness of the LASSO_Q when applied to real data. In this study, patients were randomized either to 2 months of closed-loop therapy from dinner to waking up, plus open-loop therapy during the day, or, to 2 months of all-day open-loop therapy. Here, we selected only the data collected during the all-day open-loop phase, since we are working on this specific type of therapy setting. Only meals and postprandial intervals lasting 4 hours were considered in the analysis. Intervals containing rescue carbohydrate intakes or correction boluses were excluded, to fairly evaluate the effectiveness of the mealtime IB. Moreover, only intervals with a combination of positive preprandial ROC and postprandial hyperglycemic event occurrence (scenario A) and of negative preprandial ROC and postprandial hypoglycemic event occurrence (scenario B) were taken into account because, in these cases, the expected result of an effective ROC-based mealtime IB calculation is already known: an increased IB amount, in scenario A, and a decreased IB amount, in scenario B, when compared to the original dose injected by the patient. Moreover, we selected intervals with a magnitude of hypo- and hyper event greater than 10 % of the total time window, to avoid irrelevant episodes. The resulting dataset is composed of 218 glycemic traces, 169 for scenario A, and 49 for

scenario B.

Application of ReplayBG: Briefly, the physiological model of choice uses the glucose-insulin minimal model proposed by Bergman et al. [79] as the core to describe both the effect of the insulin action and the glucose rate of appearance on plasma glucose dynamics through time. However, since neither the insulin action nor the glucose rate of appearance is usually available (as in our case), the model has been expanded with the models of Schiavon et al. [80] and Dalla Man et al. [81], which allows the final, composite model to take, as inputs, (available) exogenous insulin infusion and meal carbohydrate intakes. Further details on the model are reported in Appendix A.

Since the objective was to describe the effect of carbohydrate intakes and insulin bolus on glucose concentration, the physiological model was identified separately for each of the 218 traces by adopting the Bayesian framework, described in [77], which makes it possible to, effectively, circumvent any undesired non-identifiability issues and, also, provides point estimates of unknown model parameters by exploiting a Markov-Chain Monte Carlo strategy. Other details regarding both the physiological model and its identification procedure can be found in Appendix A. Consequently, we obtained 218 different parameter sets (PS), one for each trace. We then analyzed, trace-by-trace, the results obtained from this identification procedure. In particular, we decided to discard those traces whose pre-prandial BG value was outside the euglycemic range from the final evaluation for two main reasons: (i) they do not belong to our bolus calculator domain of validity, since it was trained only on glucose traces with this initial condition; (ii) we observed that glucose traces with pre-prandial BG concentrations outside the euglycemic range resulted in model parameters which were not physiologically plausible. Therefore, the resulting dataset is composed of 129 glycemic traces, 110 in scenario A and 19 in scenario B.

Then, to quantify, for each of the 129 traces, the glycemic outcomes resulting from the use of LASSO_Q, we set up a 4-hour long scenario where we simulated the corresponding trace five times using the real insulin dose input and the doses computed with LASSO_Q, BU, SC, ZI and SF respectively. For each simulation, we quantified glycemic control in terms of TIR, TBR, and TAR.

4.4.3 Results

The performance of the LASSO_Q model was assessed separately in scenario A and scenario B, as previously outlined. The times in each glycemic range are reported either as mean ($\pm$ standard deviation) for Gaussian distributed metrics and median [interquartile range] or otherwise. For this purpose, the non-Gaussian nature of each distribution was checked using the Lilliefors test with a 1% confidence level. The resulting metrics are reported in Table 4.5 for scenario A, and Table 4.6 for scenario B.

To compare the results obtained through the adoption of LASSO_Q in scenario A versus the other methods considered, we reduced TAR while maintaining comparable results in terms of TIR and without inducing hypoglycemia. However, in scenario B, LASSO_Q considerably increased the TIR and reduced the TBR. Indeed, while the LASSO_Q median value of 20.83% was equal to the value obtained with the administered bolus, both the 25^{th} and the 75^{th} percentiles were about 5% lower. This same result was also observed when LASSO_Q was compared with the other methods considered. In conclusion, the application of the proposed LASSO_Q model to the real data supports the positive results obtained in silico.

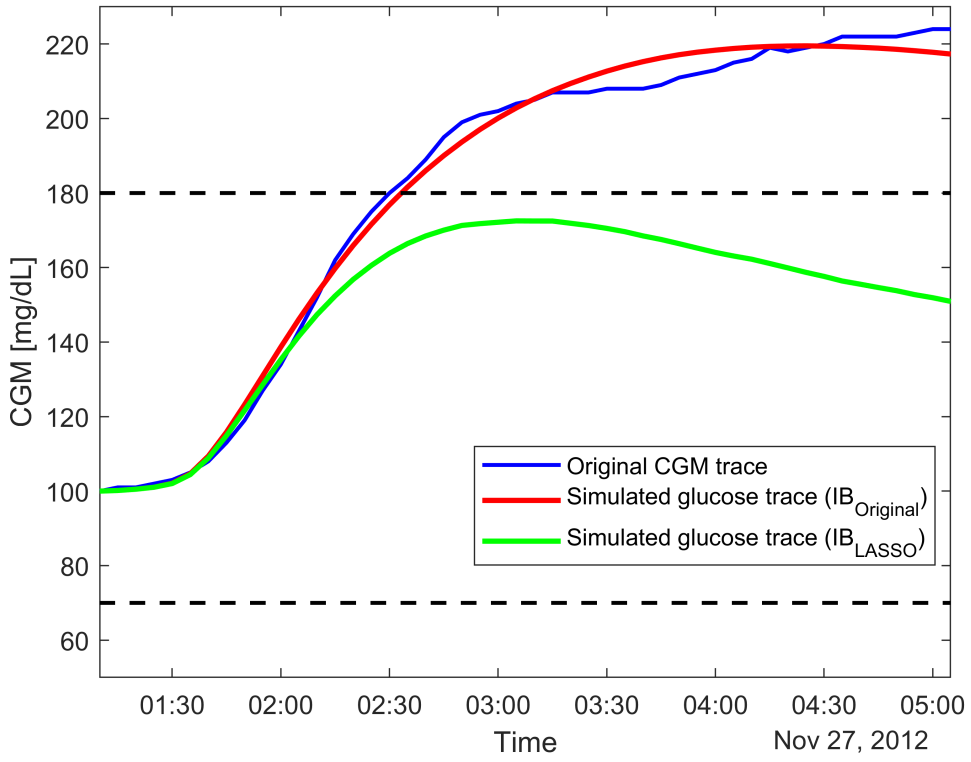

Figure 4.5: Representative example reporting the original CGM data (in blue), together with the simulated glucose trace obtained using $IB_{Original}$ (in red) and the glucose trace obtained using IB_{LASSO} (in green). The proposed model could have prevented the hyperglycemic event by increasing the dose.

4.5 Summary of the main findings and ideas for new insulin dosing strategies

By adopting supervised learning techniques, we developed multiple models, both linear and nonlinear, for mealtime IB calculation, to improve the SF traditionally used for insulin dosage and, hence, improve the quality of glycemic control while increasing model complexity. We assessed the performance of these models by evaluating the goodness-of-fit (RMSE), quantifying each model's ability to approximate the optimal insulin dose (IB_{OPT}), and then compared them with commonly adopted glycemic control indices (TIR, TAR, TBR, and BGRI). We found out that all the proposed models outperformed the

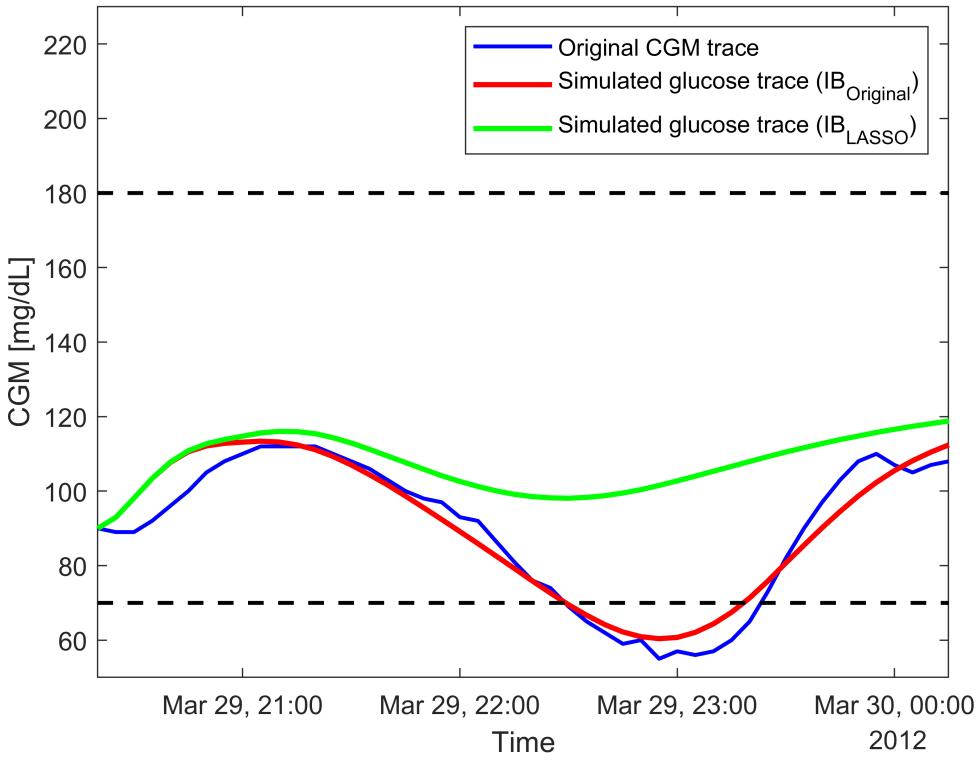

Figure 4.6: Representative example reporting the original CGM data (in blue), together with the simulated glucose trace obtained using $\text{IB}_{Original}$ (in red) and the glucose trace obtained using IB_{LASSO} (in green).The proposed model could have prevented the hypoglycemic event by decreasing the dose.

standard guidelines for insulin dosing, particularly, among the linear models proposed in Section 4.2, the in silico evaluation showed that LASSO_{Q} and LASSO_{QI} performed better than MLR and LASSO models, being able to better approach the IB_{OPT} thanks to the addition of quadratic and interaction terms between the features as input variables. At a later stage, we applied nonlinear methodologies, i.e. RF and GBT, to describe the nonlinear connection among the input variables and the output, that is the mealtime IB, to improve the linear regression model developed in 4.2. Specifically, results obtained within the simulated framework lead to statistically significant improvement in glycemic control, in terms of TBR, over the linear technique. Finally, to combine the different predictive abilities of both linear and nonlinear models, we designed

Table 4.5: Comparison of metrics obtained from scenario A, for SF, state-of-the-art methodologies and LASSO_Q. Median and interquartile ranges are reported for TBR, TAR, and mean and standard deviation for TIR.

Metric	Original bolus	LASSO_Q	SF	BU	SC	ZI
TBR %	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-0.00)
TAR %	41.67 (20.83-62.50)	38.54 (20.83-70.83)	41.67 (22.92-70.83)	39.58 (20.83-13)	40.63 (20.83-66.67)	39.58 (18.75-62.50)
TIR %	54.72 ($\pm$ 27.11)	54.45 ($\pm$ 27.88)	53.05 ($\pm$ 27.87)	54.47 ($\pm$ 28.08)	54.15 ($\pm$ 27.73)	54.47 ($\pm$ 28.07)

Table 4.6: Comparison of metrics obtained from scenario B, for SF, state-of-the-art methodologies and LASSO_Q. Median and interquartile ranges are reported for TBR, and TAR, and mean and standard deviation for TIR.

Metric	Original bolus	LASSO_Q	SF	BU	SC	ZI
TBR %	28.83 (15.10-39.98)	20.83 (10.94-29.17)	27.08 (12.50-44.27)	25.00 (12.50-44.27)	22.92 (12.50-44.27)	22.92 (9.38-44.27)
TAR %	0.00 (0.00-0.00)	0.00 (0.00-0.00)	0.00 (0.00-9.38)	0.00 (0.00-9.38)	0.00 (0.00-9.38)	0.00 (0.00-9.38)
TIR %	69.74 ($\pm$ 19.19)	72.26 ($\pm$ 20.55)	67.76 ($\pm$ 22.10)	67.65 ($\pm$ 22.15)	67.76 ($\pm$ 22.10)	67.21 ($\pm$ 22.47)

an ensemble technique based on dynamic voting. Preliminary results showed the ability of the proposed model to combine the linear and nonlinear model, i.e. LASSO and RF respectively, based on each regressor's local performance within a specific subdomain. Indeed, by testing the proposed method in silico, we found positive results in terms of the trade-off between TAR and TBR, while keeping a high median TIR. Despite the positive results obtained, the improvement in terms of glycemic control provided by an increased model complexity was not remarkable. This aspect was taken into account, together with interpretability issues, when selecting as the final model a linear one, i.e. LASSO_Q. A retrospective evaluation of this model on real data was performed by leveraging the ReplayBG tool, to evaluate the efficacy of the selected model within a scenario which includes error sources, such as CGM measurement error or carbohydrate counting error. Positive results, obtained through simulations, were confirmed by the aforementioned retrospective analysis of real data. Indeed, the application of LASSO_Q provided a reduced TAR within

scenarios including meals with postprandial hyperglycemia and a lower TIR within scenarios including meals with postprandial hypoglycemia.

Therefore, preliminary results suggest that the combination of both features on BG dynamics at mealtime and supervised learning-based models are key to improving the calculation of IB and further reducing the risk of hypo- and hyperglycemia, during and after the meal. However, the design of a supervised learning framework is far from being trivial in such a context, due to the difficulty in retrieving the optimal target of the learning task. Indeed, when considering data collected from people living with T1D, the administered mealtime insulin bolus is in most cases sub-optimal, leading to poor postprandial glycemic control. Moreover, the insulin amount estimation is highly patient-dependent, making it difficult to train a general model which is valid for different subjects. Hence, the impossibility of having reliable data, which includes an optimal insulin bolus, together with the need for a model which is tailored to the patient's needs, supports the application of reinforcement learning (RL) for such a task.

Chapter 5

Development of a reinforcement learning-based insulin bolus calculator

In this Chapter, we discuss the implementation of a double deep Q-learning (DDQ) algorithm aimed at optimizing the prandial insulin dosage. First, we will describe in detail the core algorithm adopted in our strategy, i.e., DDQ-learning, which is a variant of Q-learning, a popular reinforcement learning (RL) technique that learns the value of an action in a particular state and aims at finding an optimal policy in the sense of maximizing the expected value of the total reward over any successive step, starting from the current state, by training a rational agent [82]. Specifically, in DDQ-learning, the agent is controlled by a neural network, called deep Q network (DQN). Then, we will outline the two-step learning framework which allowed us to train the models together with the performances of the algorithm resulting from the in-silico testing.

5.1 Rationale: beyond supervised learning

As mentioned in the previous Chapter, a supervised learning framework presents several limitations in such a context. Despite many publicly available databases, which include information collected from T1D subjects, can be found, these data most often lack reliability. Indeed, patient error within T1D therapy is frequent, due to the numerous amount of tasks which should be performed by the subject and the difficulty in correctly reporting the information related to the therapy, such as the precise amount of CHO intake of the meal. Even more, a proper estimation of the insulin dose related to a meal is a challenging task, which often leads to incorrect insulin dosing and thus, suboptimal glycemic control. All these aspects lead to unreliable data, that could worsen the quality of the proposed supervised model. For this reason, we decided to abandon the idea of creating a supervised learning framework, to overcome the aforementioned issues, by employing RL, which does not need labelled data to train the model, since it relies on a trial-and-error process.

The work of Zhu et al. [61] proposed a preliminary study in this direction, by applying a deep RL algorithm (deep deterministic policy gradient) for the design of an insulin bolus advisor. The use of this methodology to develop a bolus calculator showed encouraging results, suggesting that RL is suitable for such a task. However, the simulated environment used in [61] did not consider relevant variability sources which could impact glycemic control, e.g., CGM measurement error, and the size of virtual cohort is relatively small (n $= 20$).

In this study, we extended the work proposed in [61] by applying a personalized and adaptive DDQ-learning algorithm within a highly realistic simulation scenario, which includes multiple error and variability sources, as highlighted in the following sections.

5.2 Double deep Q-learning

Hereafter, we will formulate the insulin dosing task as an RL problem. First, in Section 5.2.1 we will describe in detail the core algorithm adopted in our strategy, i.e., DDQ-learning, which is a variant of Q-learning, a popular RL technique that learns the value of an action in a particular state and aims at finding an optimal policy in the sense of maximizing the expected value of the total reward over any successive step, starting from the current state [82]. Then, in Section 5.2.2, we will present how we integrated DDQ-learning for the specific purpose of developing a new mealtime insulin bolus calculator.

5.2.1 Background on double-deep Q-learning

In general, the goal of RL is training an agent to perform a task thanks to the interaction with a defined environment, by assigning a specific reward to each agent's action [83]. In particular, for each discrete time step t the environment can be represented by the state vector (s_t), which better describes the current status of the environment among all possible vectors of the state space S. Hence, the state at time t, is used by the agent to choose an action (a_t) from all the possible sets of actions (A), according to its policy ($\pi : S \rightarrow A$). The action, which is selected based on the policy π, is applied to the environment, that evolves into the subsequent state (s_{t+1}). Finally, the action related to the specific state is evaluated through a reward (r_t), that is assigned to the corresponding state-action pair (s_t, a_t).

In this context, the goal of the algorithm is learning an optimal policy π which maximizes the cumulative discounted future reward (i.e., the return G_t):

$$G_t = \sum_{\tau=t}^{\infty} \gamma^{\tau-t} r_{\tau+1} \tag{5.1}$$

where γ is the discount factor, which takes values between [0,1], and determines how much the rewards in the distant future should be considered within the calculation.

In Q-learning, in order to learn the optimal policy and maximize the reward defined in Equation (5.1), the so-called action-value function is used [82]:

$$Q_\pi(s_t, a_t) = \mathbb{E}[G_t | s_t, a_t, \pi] \tag{5.2}$$

which is the expected return given the state s_t and action a_t under a specific policy π, representing the goodness of the action for the given state. The optimal action-value function $Q^*(s, a)$ can be defined as the one maximizing Equation (5.2), and satisfying the Bellman optimality equation [82]:

$$Q^*(s_t, a_t) = \mathbb{E}[r_{t+1} + \gamma \max_{a_{t+1}} Q^*(s_{t+1}, a_{t+1}) | s_t, a_t] \tag{5.3}$$

When the number of possible state-action pairs is high, the amount of time required to explore each state and apply a specific action is impracticable. Thus, generating the so-called Q-table, which stores $Q^*(s_t, a_t)$ for each state-action pair, becomes an excessively time-consuming process. To overcome this issue, in deep Q learning, the exact action-value function is replaced by a function approximator based on deep learning, i.e., a deep Q network (DQN). A DQN is a deep multi-layered neural network parametrized by its weight vector θ, which for a given state s provides as output the approximation of $Q^*(s, a)$, i.e., $Q_\theta(s, a)$ for all the possible actions in A.

Improving learning stability of DQN

In this work, we took advantage of two important factors related to the DQN algorithm, which increase learning stability: the experience replay and the tar-

get network [84]. By leveraging experience replay, we stored the agent's experiences in the form $(s_t, a_t, r_{t+1}, s_{t+1})$ in a cyclic buffer called replay memory, from which a minibatch was randomly drawn to update the DQN, instead of performing such a task at each time step and thus leading to poor stability. The use of experience replay also helps in breaking the correlation between two consecutive samples while training the network. The second stabilizing method was obtained by means of a target network, which is a separate network having weights θ^-, that is initially equal to the ones of the network enacting the policy. During training, the weights of the target network θ^- were updated to match the policy network θ after a fixed number of steps [84].

In this study, we applied Double Deep Q-learning (DDQ), a variant of deep Q-learning, which leverages the target network θ^- to tackle the maximization bias issue, i.e., the systematic overestimation of the action-value function due to the maximization step in Equation (5.3), which characterized such an algorithm. The max operator in DQN, in Equation (5.3), uses the same values both to select and evaluate action. This makes it more likely to select overestimated values, resulting in over-optimistic value estimates. To prevent this, DDQ-learning decouples the action selected from the action evaluation by leveraging the target network parametrized by θ^- [85]. Hence, in DDQ-learning the target Q-value is computed as follows:

$$Q(s_t, a_t) = r_{t+1} + \gamma Q_\theta(s_{t+1}, \underset{a_{t+1}}{\operatorname{argmax}}\, Q_{\theta^-}(s_{t+1}, a_{t+1})) \tag{5.4}$$

Note that, the selection of the action is due to the policy network θ, while the second network θ^- is used to fairly evaluate the value of this policy [85].

Fig. 5.1 visually summarizes the different steps and elements which compose the DDQ-learning algorithm.

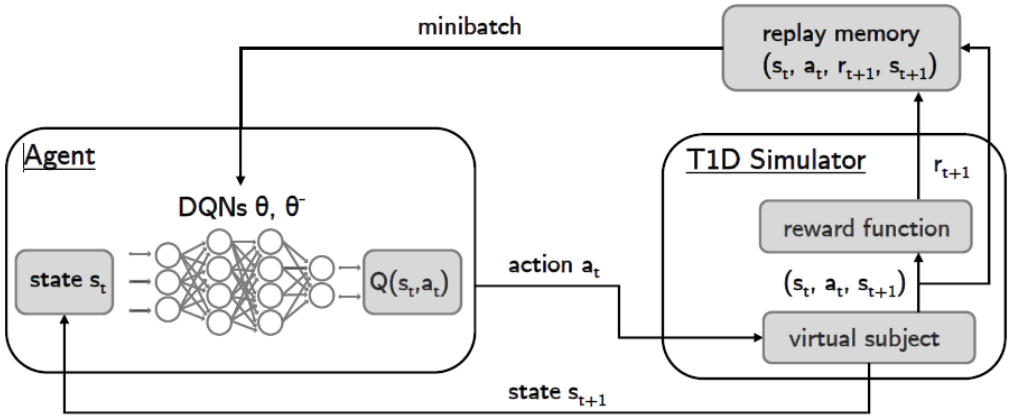

Figure 5.1: Representative scheme of the DDQ-learning framework applied to the T1D simulation environment. Within the **Agent** block, the state s_t feeds both the policy and the target network, which will be used to approximate the $Q(s_t, a_t)$. Then, the action a_t associated with the maximum estimated Q-value will be an input of the **T1D Simulator** block, and in particular to the virtual subject, which will receive the insulin dose corresponding to the chosen action. At time step $t + 1$ the environment evolves into the subsequent state s_{t+1}, hence the reward r_{t+1} is computed and the transitions $(s_t, a_t, r_{t+1}, s_{t+1})$ are stored within the replay memory and will be used every N steps to update the policy network.

5.2.2 Development of the insulin bolus calculator based on double deep Q-learning

In this section, we present how we used the DDQ-learning algorithm to develop a mealtime insulin bolus calculator. In particular, at mealtime t, the patient needs to estimate and inject the insulin dose to counteract the glycemic excursion due to the meal. Within this framework, the mealtime condition is represented by a specific state s_t, composed of easily accessible physiological parameters of the T1D individual as described in 5.2.2. Based on the current mealtime status, the best action a_t selected by the policy was applied to the patient, i.e., the best insulin amount to be injected, as defined in 5.2.2. Lastly, after the insulin amount delivery, the postprandial glycemic outcome was evaluated through the reward function specified in 5.2.2.

State vector choice

In this work, we used easily accessible variables together with patient-specific therapy parameters to describe the mealtime condition, which provides a comprehensive view of the prandial status.

We defined the state vector at time t (s_t) by extracting, for each meal, the following variables: the current CGM measurement G [mg/dL]; the carbohydrate content of the meal (CHO) [g]; the CGM rate of change (ROC) [mg/dL/min], which gives insight into the glucose dynamics, by indicating whether glucose level is falling or rising and to what extent; the prandial insulin-to-carbohydrates ratio (CR) [g/U], a therapy parameters indicating how many grams of CHO will be covered by one unit of insulin and which could vary from meal to meal; together with the prandial insulin dose computed through the standard therapy (BC_S) [U] as in Equation 1.1. In this process of feature selection, we followed our previous work's rationale [46], which involved the exclusion of variables that, in this specific application, assumed a constant value for each meal, e.g., the body weight (BW), the correction factor (CF), the target glucose (G_t). As a result, we defined s_t as:

$$s_t = \{G_t, CHO_t, ROC_t, CR_t, BC_{S_t}\}.$$

Set of possible actions

In our scenario, there are multiple types of actions which could potentially be adopted, such as the estimation of the insulin dose itself, or the correction of the dose suggested by BC_s. Since the number of possible insulin amounts associated with a state is considerably high, we decided not to directly estimate the mealtime bolus, but to correct the dose suggested by BC_s, which thus will be considered as a starting value. This choice was made to avoid an exces-

sive number of possible actions, which could lower the performance of our algorithm. Hence, we defined an action a_t as a percentage modulation of the mealtime insulin dose suggested by the standard therapy. In particular, the BC_s dose can be decreased or increased by a percentage α, which ranges from these possible values $\alpha = \{\pm 25, \pm 20, \pm 10, 0\}\%$. The motivation behind this choice lies in the results obtained in our previous work [46] where multiple state-of-art insulin adjustment methods were preliminarily assessed in silico. The literature approach which was correcting the dose based on a percentage modulation resulted in the safest method under specific mealtime conditions, potentially indicating that the insulin correction should also depend on the entity of the meal and not be fixed to constant values.

Following this rationale, at each time step t, the chosen action $a_t = \alpha_t$ was applied to the bolus calculator as follows:

$$BC_{ddqn}(t) = BC_s(t) + \alpha_t BC_s(t) \tag{5.5}$$

where $BC_{ddqn}(t)$ represents the insulin amount suggested by the DDQN algorithm at time step t.

Reward function

Among all the possible reward functions which could be employed to evaluate postprandial glycemic control, we followed the structure of the one presented in the preliminary work of [61], which showed promising results for such a purpose. The aim being evaluating the glycemic control following the current meal at time instant t, we assigned a weight to each postprandial CGM sample between the current time instant t and t^*, which corresponds either to the following meal or a postprandial 6 hours interval if the next meal occurred at a longer time distance. In particular, the reward equation is reported in Equation

(5.6):

$$r_t = \frac{1}{t^* - t} \sum_{k=t}^{t^*} f_R(G_k) \tag{5.6}$$

where $t^* = min(t + 6h, t + 1)$. Note that weights were assigned differently based on the glycemic range. Particularly, hypoglycemia was penalized more compared to hyperglycemia, since a hypoglycemic excursion is riskier than a hyperglycemic excursion having the same amplitude. The selected rewards were tuned by trial-and-error and associated with each glycemic interval as follows:

$$f_R(G_k) = \begin{cases} +0.5 & \text{if } 70 \leq G_k \leq 180 \\ -0.9 & \text{if } 180 < G_k \leq 200 \\ -1.2 & \text{if } 200 < G_k \leq 250 \\ -1.5 & \text{if } 250 < G_k \leq 350 \\ -1.8 & \text{if } 30 < G_k < 70 \\ -2 & \text{else} \end{cases} \tag{5.7}$$

As reported in Equation (5.7), a glucose level within the target range was associated with a positive reward. The highest penalty was assigned to values in hypoglycemia, while hyperglycemic values were not penalised equally, having different weights based on the severity of the interval. Fig. 5.2 depicts the different weights assigned to each glycemic interval.

Implementation details on the DDQN algorithm

In this Section, implementation details are reported to further clarify the training process. During the learning phase, all the hyperparameters highlighted in Section 5.2.1 have been set, following a trial-and-error procedure, as reported in Table 5.2, while the weights of the DQNs were initially initialized as random.

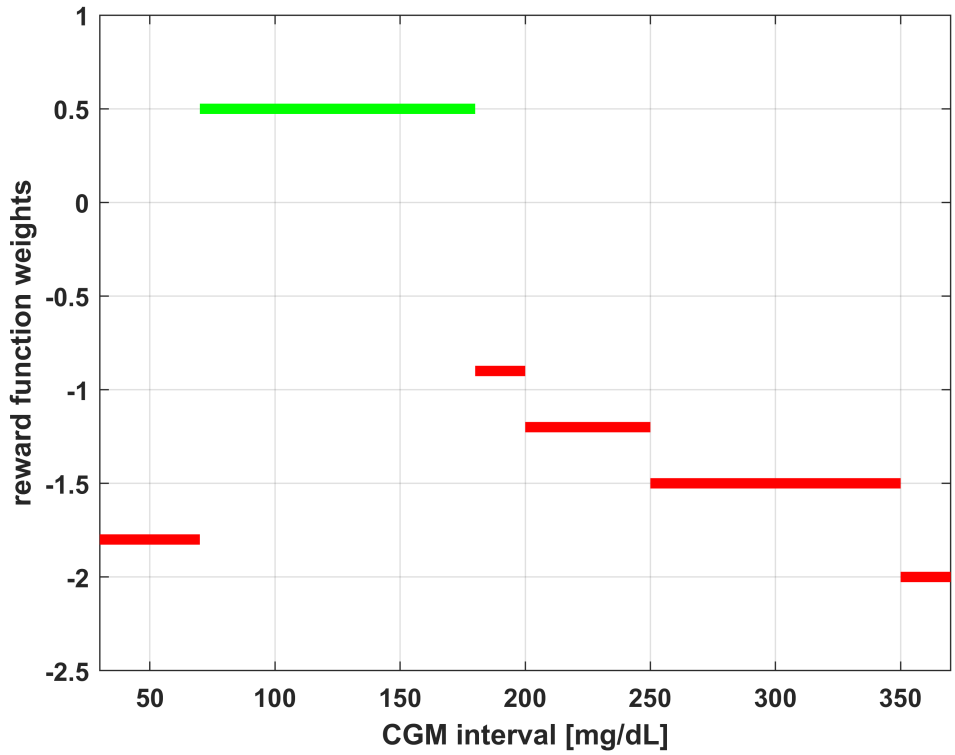

Figure 5.2: Reward function employed for the proposed DDQN algorithm. Each CGM interval is associated with a constant value used within the reward function. Green line represents the euglycemic range, while the red lines indicate intervals associated with adverse events.

Before starting DDQN training, the replay memory introduced in Section 5.2.1 was first filled by applying random actions chosen from the set of possible actions of Section 5.2.2. After this phase, characterized by actions chosen completely random, we started the learning procedure by selecting either a random action with probability ϵ or an action based on the agent with probability $1 - \epsilon$. In particular, the agent makes use of a DQN, having two hidden layers composed of 32 and 16 nodes respectively, which maps state-space ($\mathbb{R}^5$) to action-space ($\mathbb{R}^7$). The weights of the policy network were updated at the end of one episode, according to minibatch gradient descent, by randomly sampling from the replay memory, while the target network was updated after a fixed number of episodes.

5.3 Simulation framework

The proposed algorithm is trained and tested by means of the UVA/Padova T1D Simulator.

Indeed, due to the nature of RL techniques, the use of a simulation environment is particularly suitable, being the learning of the model achieved through the interaction with the environment and thus, following a trial and error procedure. Performing such a process on a virtual subject is key to avoiding dangerous situations within a real clinical setting.

In this work, we used an updated version of the FDA-accepted UVA/Padova T1D Simulator [60], which includes several elements that enhance the realism of the simulated scenarios. These improvements include a dedicated model to describe the dawn phenomenon and time-varying therapy parameters modelled as in [37], a model of CGM sensor measurement error [86, 87], and a behavioural model of people with T1D [88].

The virtual cohort was composed of 100 adult subjects, and the total pop-

ulation was divided into two different subgroups, hereafter labelled as Set A and Set B. We obtained the two subsets by randomly selecting 50 subjects for each group while performing a stratified split, in which distribution of the average CR pattern, the CF, and the body weight (BW) variables between the two groups remained the same. The stratification was performed to ensure homogeneity between the two sets, being the aforementioned parameters used to identify the subjects used to train the population models within the first step of the learning framework. Hence, Set A was employed for the development of the first learning step, while Set B was retained to fine-tune the generalized models and, lastly, test the performances. The first stage, applied on Set A, consisted of a long-term training of K population models on 1200 simulated days, while the personalization of the resulting DDQN models was carried out on Set B, on a 180-day simulation as performed in [61]. Such a long simulation time was needed since meals are infrequent events within a day (three meals per day).

One day of simulation included three meals per day: breakfast, lunch and dinner. Both meal timing and CHO content of the meals were extracted from uniform distributions to match the data reported in [89]. Table 5.1 reports the minimum and maximum values of the distributions from which the meal timing and CHO amount are drawn.

Moreover, to further improve the realism of the simulated scenarios, meal carbohydrate counting error has been modelled and added to the original meal amount as described in [90]. Of note, no corrective actions (rescue carbohydrates intakes or corrective insulin boluses) were included during the simulations to fairly compute the reward corresponding to the action chosen by the algorithm. Indeed, a corrective action could add confounding factors within the rewarding process, since it would bring glycemic levels back into the normoglycemic range, not letting us assess the efficacy of the chosen action alone.

Table 5.1: Minimum and maximum values of the possible CHO amount and time of consumption for the different meals.

Meal type	CHO amount [g]	Meal timing
breakfast	[19-97]	[6.30 am - 8.00 am]
lunch	[31-124]	[11.30 am - 1.00 pm]
dinner	[28-140]	[6.00 pm - 8.30 pm]

As a result, the only control action a_t allowed between mealtime at time t and the following meal at time $t+1$ was represented by the meal insulin bolus suggested by the policy.

5.4 Learning of the deep Q-network (DQN)

Learning of our model was performed on the first set of 50 subjects extracted in Section 5.3, following the same two-step learning rationale used in Zhu et al. [61] In particular, in the first step we used a long-term simulation to train a set of K population models able to represent the inter-subject physiological variability observed in people with T1D, as described below in Section 5.4.1 (Step 1). Secondly, for a given subject, a specific model is selected between the K population models and then personalized to better fit his/her peculiar physiology, as reported in Section 5.4.2 (Step 2) .

5.4.1 Step 1: Sub-population model training

To obtain K sub-population models which can describe the variability in T1D subjects' physiology, we performed a sub-population model training by dividing Set A of subjects into K different groups that share similar characteristics. To achieve this goal, each subject was described with three patient-specific parameters, i.e., the average CR value of the daily pattern, the CF variable, and the BW. Then, we applied the K-medoids clustering algorithm [91], to divide the 50 subjects composing the used population into K clusters. In particular, K

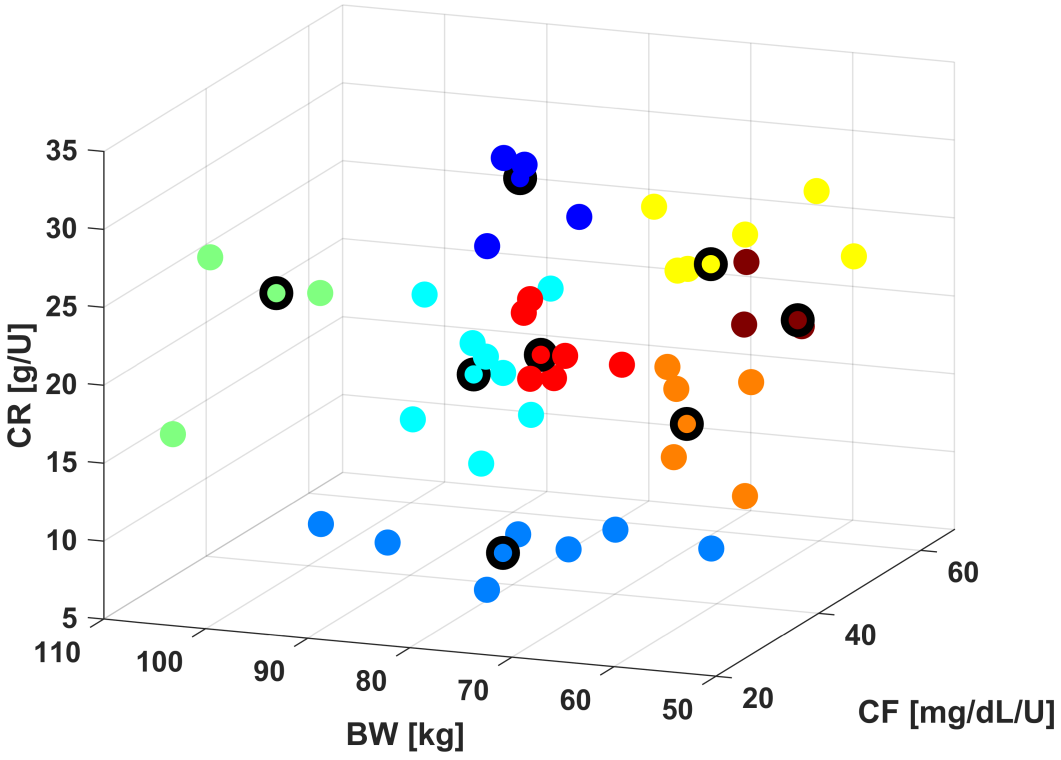

Figure 5.3: Scatter plot of the 8 different clusters resulting from the application of the K-medoids algorithm. Each of the 50 subjects is represented by a specific BW, CF and average CR value. The sub-population subjects, i.e., the medoids of each cluster, are outlined in black.

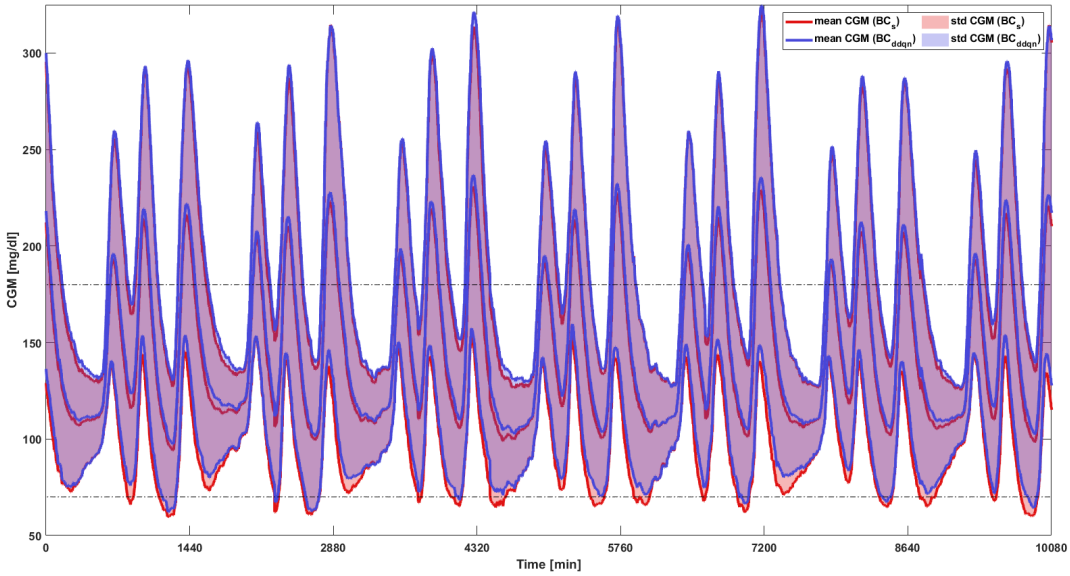

Figure 5.4: CGM mean and standard deviation intervals resulting from the virtual subjects belonging to the test set are reported for a representative one-week-long simulation. CGM values related to the standard insulin dosing and DDQN bolus calculator are shown in red and blue respectively. Dashed lines indicate the normoglycemic range.

= 8 was identified as the optimal number of clusters by employing the elbow method. Finally, we selected the K = 8 representative subjects as the medoids of each cluster. Fig. 5.3 depicts, with different colours, the 8 identified clusters and the respective representative subjects, highlighted with a black circle.

5.4.2 Step 2: Personalization of DQN-learning models on patient-specific data

In the second step, the sub-population DDQNs were fine-tuned and thus, further personalized on each subject of Set B. Indeed, after having defined the cluster to which each virtual patient of Set B belongs, the DDQN weights were initialized with the one resulting from the previous phase, i.e., θ_k and θ_k^-, where k indicates the k-th sub-population subject.

By setting specific safety constraints, which limits the meal insulin correc-

Table 5.2: Hyperparameters used for both the sub-population model training and personalized model tuning.The hyperparameters do not differ among the virtual subjects.

Hyperparameter	Value
Number of random episodes before learning starts	30
Sub-population ϵ-greedy exploration	$0.9 \rightarrow 0.1$
Number of episodes after which the target DQN is updated	50
Replay memory maximum size	800
Discount factor γ	0.95
Minibatch size	32
Learning rate	0.001
Hidden nodes of the DQNs	[32, 16]

tion, this phase of individualization could also be safely carried out on the real subject within a clinical trial setting. In addition to the possibility of introducing security constraints, this phase may be performed using simulation tools which leverage real data, such as [77], thus enabling the fine-tuning of the generalized model on simulations derived from patients' real data.

5.5 Results

The assessment of the RL algorithm was performed on 60 simulated days, where the experimental set-up was the one described in Section 5.3, and the population cohort was composed of 50 subjects (Set B). Moreover, to evaluate the efficacy of the proposed methodology, results were compared to 60 days of simulation in which the standard bolus calculator reported in Equation (1.1) was used for the insulin dosing task. It should be pointed out that, the daily scenarios in terms of mealtimes and CHO amounts composing the meal within the 60-day testing period differ from those within the 180-day personalization period of Section 5.4.2. Lastly, to ensure a fair comparison between the two insulin dosing methods and to allow replicable analysis, the seed of the random number generator was fixed for each virtual subject.

Table 5.3: Values related to median and interquartile ranges of TBR [%], TIR [%], TAR [%], N_{hypo} and N_{hyper} are reported for both the standard and DDQN bolus calculators.

	TBR [%]	TIR [%]	TAR[%]	N_{hypo}	N_{hyper}
BC_s	8.78 (2.90, 12.44)	68.35 (64.87, 72.65)	22.24 (15.89, 30.89)	1 (0, 1)	2 (2, 3)
BC_{ddqn}	4.17* (2.70, 8.17)	70.08 (66.49, 76.55)	23.47 (18.25, 28.64)	0 (0, 1)	2 (2, 3)

*Statistically significant compared to BC_s with $p < 5\%$

From the 60-day glucose profile, we extracted several metrics evaluating glycemic control, which are widely adopted by the diabetes research community [92, 93] and reported in Section 2.3.2 within Chapter 2. In particular, we derived three metrics related to the the percentage of time spent within the different glycemic ranges, that is TIR, TAR and TBR. Moreover, we computed two popular indices used to quantify the risk of hypo- and hyperglycemia, namely the low blood glucose index (LBGI) and high blood glucose index (HBGI) respectively, together with the overall BGRI, which sums the two contributions in one risk index, indicating the goodness of the overall glycemic control [66]. In addition to these metrics, the median number of hypo- and hyperglycemic events per day was extracted to assess the benefit introduced by the proposed algorithm. Lastly, to evaluate the statistical significance of the resulting metric distributions we applied a paired t-test with significance level equal to 5% to those metrics having a Gaussian distribution based on the Lilliefors test (i.e., TIR, TAR) [94], while the Wilcoxon test with significance level equal to 5% was used for the TBR metric, which showed a non-Gaussian distribution.

In Fig. 5.4, the mean and the corresponding standard deviation extracted from all the glucose profiles of the virtual subjects belonging to the testing set are shown for both the DDQN algorithm and the standard therapy. It is noted that, considering the negative shift of the standard deviation, in most hypoglycemic events, the nadir is considerably reduced when BC_{ddqn} is used, while the increase in terms of maximum hyperglycemic value introduced by

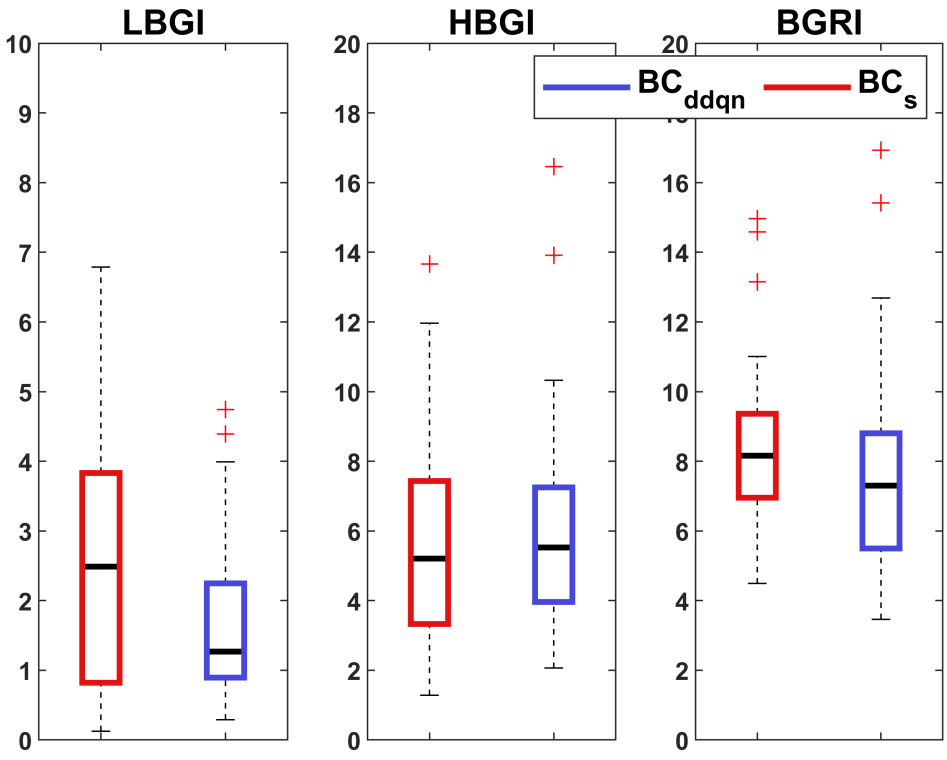

Figure 5.5: Distributions of LBGI, HBGI, and BGRI resulting from the testing phase are reported in blue for the DDQN bolus calculator, while in red for the standard bolus calculator.

the DDQN algorithm is almost negligible.

The aforementioned remarks are consistent with the results obtained by analyzing the metrics related to the time in different glycemic ranges, as reported in Table 5.3. Indeed, the advantage brought in terms of hypoglycemia reduction is reflected in a significantly lower distribution of TBR for BC_{ddqn} compared to BC_s, with a median value of 4.17% and 8.78% respectively. This was achieved without negatively impacting on hyperglycemia, as can be seen from the TAR distributions in Table 5.3, which reports a median value of 23.47% and 22.24% for the DDQN and standard therapy respectively, not showing a statistical significance between the two distributions. In general, the benefit introduced by the proposed algorithm is positive, since both the interquartile and the median value related to the TIR are improved.

The positive impact given by the DDQN method was also confirmed when observing the median number of adverse events per day, shown in Table 5.3. Indeed, median N_{hypo} is reduced from 1 event per day to zero, while N_{hyper} remains unchanged for both methods.

To ensure that the moderate increase in TAR is not influencing negatively the overall glycemic control, the risk metrics described previously were analysed and reported in Fig. 5.5. As expected, the LBGI is significantly reduced, while HBGI showed a slight increase in the median value. However, the BGRI, which summarizes the two aforementioned metrics, showed an improvement, by decreasing the median value from 8.2 of the benchmark to 7.3 of the BC_{ddqn}. Also, the total amount of insulin injected for each subject remained stable for both methods, suggesting that the significant improvement in terms of metrics related to hypoglycemia is not simply due to a decrease in insulin dosage, but to a more effective redistribution of such hormone.

The reported results pointed out that, in general, the use of BC_{ddqn} could provide a beneficial impact on glycemic control, by considerably reducing not

only the occurrence but also the duration of hypoglycemia, without significantly affecting hyperglycemia.

5.6 Summary of the main findings and possible future developments

In this work, we proposed a mealtime insulin bolus calculator based on a DDQN algorithm, which leverages a high level of personalization and adaptation through a two-step learning framework combined with a clustering procedure. This method allowed us to fine-tune the generalized model related to the sub-population subject which shows more similarity to the patient. Such individualization of the insulin dosing was implemented to deal with the multiple variability sources introduced by the updated and highly realistic version of the UVA/Padova T1D Simulator used in this work, as described in Section 5.3.

The presented bolus calculator was tested within the simulated environment for a 60 days simulation, and compared to the state-of-art method for insulin dose calculation, i.e., the standard bolus calculator in Equation (1.1), by extracting different metrics widely used by the diabetes research community to assess the quality of glycemic control. Despite the challenging scenario provided by the employed simulated environment, results in terms of time spent within the different glycemic ranges, the number of adverse events and glycemic risk metrics was encouraging, showing the ability of the proposed method to completely avoid or mitigate the magnitude of postprandial hypoglycemic events, while moderately improving the time in range and maintaining stable the time in hyperglycemia compared to the benchmark.

The obtained performances indicated the efficacy of the algorithm in ad-

justing the standard dosage based on the mealtime state of the subject. Moreover, the total amount of daily insulin delivered to the subject remained equal both for the standard and the proposed method, pointing out that the reduction of hypoglycemia is not due to a simple decrease of the mealtime insulin amount, but rather the algorithm is redistributing the daily amount of insulin more effectively.

Limitations of the work include the need for a relatively long period to personalize the population model, being the meal an infrequent event within a day. A possible solution to overcome this limitation may come from the use of simulation tools, such as [77], which allows to identify a model on a patient's glucose trace and to replay the scenario by changing the inputs of the model (e.g, insulin bolus), thus evaluating the effectiveness of a specific therapy.

Hence, future works will involve an assessment in a clinical setting together with the exploration of the aforementioned methodology to speed up the personalization phase. Moreover, further developments of the method may include the exploration of a simulated scenario which also takes into account the insulin sensitivity variability, a challenging factor impacting T1D management. In conclusion, the application of a DDQ-learning algorithm combined with an effective model personalization procedure allowed us to achieve promising results within the updated version of the FDA-accepted UVA/Padova simulated environment when applied for prandial insulin dose adjustment.

Chapter 6

Conclusion and future developments

6.1 Summary of thesis contributions

In general, the objective of this thesis is rising awareness of the importance of an accurate and precise prandial insulin dosage within T1D management, which is the major obstacle to optimal glucose control. As previously highlighted, we proposed different possible solutions, to demonstrate that there could be a margin for improvement over the standard guidelines and that this topic needs to be extensively addressed by the diabetes research community. Particularly, we presented different approaches based on machine learning, i.e., supervised and reinforcement learning, to improve the mealtime insulin dose estimation, thanks to the use of information on CGM trend and personalized parameters related to the patient's prandial status. In each phase of this work, we pointed out the limitations which characterize the standard guidelines for insulin dosing, highlighting the possibility of achieving an improved glycemic outcome when adopting our proposed strategies. The study reported in Chapter 2 allowed us to find out the open issues related not only to the stan-

dard guidelines, but also to the empirical methods proposed in the literature to adjust such a dosage. Indeed, by performing an ISCT within single-meal and noise-free scenarios, we compared and assessed the current approaches and verified their limitations in terms of effectiveness on glucose regulations. This work set the basis for the investigation of different data-driven strategies aimed at optimizing the prandial insulin bolus. Therefore, the following phase, reported in Chapters 3 and 4, consisted in designing novel techniques for insulin dosing based on a supervised learning framework and leveraging a simulated dataset, consisting of 100 virtual subjects, where each of them underwent different mealtime scenarios. The proposed population models, i.e., linear, nonlinear and ensemble models, led to encouraging results when tested within the simulated environment. However, the increase in model complexity was not matched with significantly higher performance, indicating that population models are not sufficiently tailored to the single patient, despite the inclusion of personalized features within the dataset. Moreover, as previously stressed within this thesis, designing a supervised learning framework for such a purpose is a challenging task, due to the possible unreliability of real data and the need of leveraging a simulated dataset. Hence, to overcome the aforementioned limitations, we extended the previous work by applying an RL algorithm, i.e. double deep Q-learning. Training the model through the interaction with the simulated environment resulted particularly suitable for such a purpose, avoiding the need for labelled data. Results obtained during the testing of the proposed algorithm were promising, showing the beneficial impact of an insulin dosing strategy which is highly personalized and promoting the investigation of future developments on this topic.

6.2 Future challenges and perspectives

Future challenges related to this work include further assessment of the proposed supervised models using a simulated environment that takes into account additional error sources. Indeed, while the models were developed within an error-free environment to eliminate any confounding factors that could have influenced the outcome of the study, the integration of the new model within a multi-meal scenario [95], including additional error sources, such as errors in patient behaviour [96], carbohydrates miscalculations [90] and sensor readings [97] [98] could be beneficial to compare the outcome of the supervised models with the results obtained through the RL-based algorithm, that was developed in such challenging framework. However, a preliminary investigation of the performance related to the supervised models using real data, which includes multiple error sources, was already performed, confirming the positive results obtained in silico.

In addition, a relevant step for the extension of this study, will involve an assessment of the presented models within a real-life setting, to evaluate the effectiveness of the machine-learning-based techniques in presence of different endogenous and exogenous factors, e.g., hormonal changes or physical activity, which are not modelled within the simulated environment.

6.2.1 Application of the presented work: integration within a decision support system

In the past decade, diabetes management has been transformed by the advent of CGM sensors. However, even though this technology provides an increased amount of information, current standard treatments are unable to leverage the information generated. The increased application of data-driven methodologies, which are capable of handling and processing this information, has led

to the development of decision-making tools and applications which can enhance the management of diabetes. In particular, decision support systems (DSSs) give the possibility to support patients with proactive and personalized decisions in any scenario of their daily living and allow them to react at a shorter time scale. Over the past few years, DSSs for chronic diseases such as diabetes have been an emerging concept in healthcare. One of the main advantages of DSS is that data can be automatically collected, transmitted, aggregated with other physiological signals, analyzed, stored and presented to the patient [99, 100, 101]. In addition, by integrating e-health and telemonitoring systems, DSSs for T1D have the potential to improve glycemic outcomes thanks to the prevention of hypo- or hyperglycemic events, reducing uncertainty when making critical self-management decisions [102]. A DSS for diabetes treatment can provide an alternative to much more complicated closed-loop systems, such as the artificial pancreas. Indeed, a wide range of users does not feel confident with the use of this device, being concerned about errors occurring in the insulin pump, and prefers an open-loop therapy, which can be greatly facilitated by the DSS. For this reason, many DSS aimed at supporting the T1D individual were proposed in the literature recently [103, 104, 105]. Most of the DSSs already available in the literature are composed of a predictive glucose module (which alerts the user whenever its BG is predicted to fall outside the safe range in the next future), an insulin suspension module (which temporarily suspends basal insulin delivery to avoid hypoglycemia when BG is critically low in a patient using insulin pumps), and a bolus calculator module to compute insulin dose at mealtime. Therefore, a straightforward application of the work presented in this thesis consists in designing a new bolus calculator module that implements the proposed methodologies to develop a new, more sophisticated, DSS that can achieve better glycemic control when compared to other state-of-the-art algorithms.

In such a context, the integration of the algorithms proposed in this thesis would support individuals affected by T1D in computing an effective insulin dosage. Following this rationale, in our recent work, we presented a preliminary DSS which integrates different modules, among which an insulin bolus advisor, to assist the patient during the daily decision-making process [106].

However, despite the increased investigation of DSSs for T1D management, many of the insulin-dosing DSS algorithms and glucose prediction DSS algorithms published in the literature included only an in silico evaluation, making it challenging to assess how well the systems will perform in a real-life setting [105]. While the simulated environments have been crucial and groundbreaking in the design and preliminary evaluation of novel treatment strategies, real-life testing should be taken into consideration to bridge the gap which is present in the literature.

Appendix A

ReplayBG

Here we discuss the nonlinear physiological model of glucose-insulin used by the ReplayBG tool. As previously done by Cappon et al. [107], the model was built starting from the physiological model available in the UVA/Padova T1D Simulator (T1DS) [37]. Moreover, as in [107], in order to permit model identification at the individual level using carbohydrate intake $CHO(t)$, exogenous insulin $I(t)$, and CGM data, the original T1DS model has been simplified reducing the number of parameters to be identified, but paying attention in maintaining its ability to effectively describe glucose-insulin dynamics. Note that CGM, CHO, and insulin data should be collected in parallel. As documented in the following, the model is composed of three main subsystems: subcutaneous insulin absorption; oral glucose absorption; glucose-insulin kinetics.

A.1 Model structure

A.1.1 Subcutaneous insulin absorption subsystem

The model of the subcutaneous insulin absorption system is a simplified version of the model incorporated in T1DS [108]. The model is composed of three

compartments and describes the absorption dynamics of exogenous insulin infusion to the plasma. Exogenous insulin $I(t)$ is infused to the first compartment, which represents insulin in a non-monomeric state. Then, "non-monomeric" insulin diffuses to the second compartment, representing insulin in a monomeric state, and eventually reaches plasma. Model equations are:

$$\begin{cases} \dot{I}_{sc1}(t) = -k_d \cdot I_{sc1}(t) + I(t-\beta)/V_I \\ \dot{I}_{sc2}(t) = k_d \cdot I_{sc1}(t) - k_{a2} \cdot I_{sc2}(t) \\ \dot{I}_p(t) = k_{a2} \cdot I_{sc2} - k_e \cdot I_p(t) \end{cases} \tag{A.1}$$

where I_{sc1} (mU/kg) and I_{sc2} (mU/kg) represent the insulin in a non-monomeric and monomeric state, respectively; I_p (mU/l) is the plasma insulin concentration; k_d (min^{-1}) is the rate constant of diffusion from the first to the second compartment; k_{a2} (min^{-1}) is the rate constant of subcutaneous insulin absorption from the second compartment to the plasma; k_e (min^{-1}) is the fractional clearance rate; V_I (l/kg) is the volume of insulin distribution; β (min) is the delay in the appearance of insulin in the first compartment.

A.1.2 Oral glucose absorption subsystem

The model of the oral glucose absorption system, taken from [109], describes the gastro-intestinal tract as a three-compartment system: the first two compartments quantify the glucose in the stomach, while the third compartment models the upper small intestine where CHO is absorbed. Model equations are:

$$\begin{cases} \dot{Q}_{sto1}(t) = -k_{empt} \cdot Q_{sto1}(t) + CHO(t) \\ \dot{Q}_{sto2}(t) = k_{empt} \cdot Q_{sto1}(t) - k_{empt} \cdot Q_{sto2}(t) \\ \dot{Q}_{gut}(t) = k_{empt} \cdot Q_{sto2}(t) - k_{abs} \cdot Q_{gut}(t) \end{cases} \tag{A.2}$$

where Q_{sto1} (mg / kg) and Q_{sto2} (mg / kg) are the amounts of glucose in the stomach in solid and liquid state, respectively; Q_{gut} (mg/kg) is the glucose concentration in the intestine; k_{empt} (min^{-1}) is the rate constant of gastric emptying; k_{abs} (min^{-1}) is the rate constant of intestinal absorption; CHO (mg/kg/min) is the ingested carbohydrate rate. Model (A.2) allows to estimate the rate of glucose appearance in plasma Ra (mg/kg/min) as:

$$Ra(t) = f \cdot k_{abs} \cdot Q_{gut}(t) \tag{A.3}$$

where f (dimensionless) is the fraction of the intestinal content absorbed in the plasma.

A.1.3 Glucose-insulin kinetics subsystem

The core model is based on a modified version of the glucose-insulin kinetics minimal model introduced in [110][111]. The model is composed of three compartments, the first describing the effect of insulin action and oral glucose rate of appearance on plasma glucose concentration, the second quantifies the impact of plasma insulin concentration on insulin action, and the last one represents the transport of glucose from plasma to the interstitium. Model equations are:

$$\begin{cases} \dot{G}(t) = -[SG + \rho(G)X(t)] \cdot G(t) + SG \cdot G_b + Ra(t)/V_G \\ \dot{X}(t) = -p_2 \cdot [X(t) - SI \cdot (I_p(t) - I_{pb})] \\ \dot{IG}(t) = -\frac{1}{\alpha}(IG(t) - G(t)) \end{cases} \tag{A.4}$$

where G (mg/dl) is the plasma glucose concentration, X (min^{-1}) is the insulin action on glucose disposal and production; IG (mg/dl) is the interstitial glucose concentration; SG (min^{-1}) is the glucose effectiveness that describes the

glucose ability to promote glucose disposal and inhibit glucose production; G_b (mg/dl) is the basal glucose concentration in the plasma; V_G (dl/kg) is the volume of glucose distribution; p_2 (min^{-1}) is the rate constant of insulin action dynamics; SI (ml/μU·min) is the insulin sensitivity; I_{pb} (mU/l) is the basal insulin concentration in the plasma; α (min) is the delay between plasmatic and IG compartments; and $\rho(G)$, is a deterministic function, introduced by Dalla Man et al. [112], that allows to better represent glucose dynamics in the hypoglycemic range by increasing insulin action when glucose decreases below a certain threshold:

$$\rho(G) = \begin{cases} 1 & \text{if } G \geq G_b \\ 1+10r_1\{[ln(G)]^{r_2} - [ln(G_b)]^{r_2}\}^2 & \\ & \text{if } G_{th} < G < G_b \\ 1+10r_1\{[ln(G_{th})]^{r_2} - [ln(G_b)]^{r_2}\}^2 & \text{if } G \leq G_{th} \end{cases} \tag{A.5}$$

where G_{th} is the hypoglycemic threshold (set to 60 mg/dl) and r_1 (dimensionless) and r_2 (dimensionless) are two model parameters, without direct physiological interpretation, set to 1.44 and 0.81, respectively.

Note that in the model of Eq. (A.4) it is assumed, without any loss of generality, the gain between $G(t)$ and $IG(t)$ to be equal to one [113] [114]. With this assumption, in practice, $IG(t)$ represents a noise-free version of the CGM data, allowing us to fit it versus $CGM(t)$ in the identification step described in the following section.

A.2 Model identification

For each model subsystem, the number of parameters to be identified was reduced as much as possible, i.e., retaining only the parameters that have the

greatest impact on model output, while setting the other parameters to population values. This selection step facilitates the identifiability of the model maintaining capability of describing the glucose dynamics observed in the data. In particular, 6 parameters were fixed to population values, i.e., V_I = 0.126 l/kg, k_e = 0.127 min^{-1}, β = 8 min, f = 0.9, V_G = 1.45 dl/kg, and α = 7 min [108] [109] [110]. The remaining parameters in eqs.(A.1-A.4) must be derived, at the individual level, from the data. Therefore, the vector $\boldsymbol{\theta}$ of the unknown parameters is composed by a total of eight variables, two related to the subcutaneous insulin subsystem, i.e., k_{a2} and k_d, two associated to the oral glucose absorption subsystem, i.e., k_{empt} and k_{abs}, and four related to the glucose-insulin kinetics subsystem, i.e, S_G, S_I, p_2, and G_b. Formally, the parameter estimation problem can be stated as the problem of determining $\boldsymbol{\theta}$ from the equations:

$$\begin{cases} \dot{\boldsymbol{x}}(t) = \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}, t, \boldsymbol{\theta}) \\ y(t) = IG(t) \end{cases} \tag{A.6}$$

where $\boldsymbol{x}(t)$ is the state vector defined as

$$\boldsymbol{x}(t) := [I_{sc1}, I_{sc2}, I_p, Q_{sto1}, Q_{sto2}, Q_{gut}, G, X, IG]^T;$$

$\boldsymbol{u}(t) := [I(t), CHO(t)]$ is the input vector; $\boldsymbol{f}(\cdot)$ is the state update function combining (A.1) (A.2), and (A.4). $\boldsymbol{f}$ depends on the set of unknown parameters $\boldsymbol{\theta}$.

As documented in the literature, (A.1), (A.2), and (A.4) are *a priori* non identifiable [108][109][111]. As such, the identification of the vector $\boldsymbol{\theta}$ at the individual level from insulin and carbohydrate intake and CGM is not trivial, since the resulting model (A.6) is not *a priori* identifiable as well.

To mitigate *a priori* non-identifiability, the Bayesian approach of Pillonetto et al. [115], i.e., Markov Chain Monte Carlo (MCMC) was adopted. In fact,

MCMC allows overcoming improper model parameters estimates that are usually achieved with maximum-likelihood-based approaches. In details MCMC, allows to obtain a point estimate of $\boldsymbol{\theta}$, $\hat{\boldsymbol{\theta}}$, by performing Monte Carlo integration over a set of N samples $\boldsymbol{\theta}_i, i = 1, \ldots, N$ generated from the posterior distribution

$$p_{\boldsymbol{\theta}|Y,U}(\boldsymbol{\theta}|Y,U) = \frac{p_{Y|\boldsymbol{\theta},U}(Y|\boldsymbol{\theta},U)p_{\boldsymbol{\theta}}(\boldsymbol{\theta})}{\int p_{Y|\boldsymbol{\theta},U}(Y|\boldsymbol{\theta},U)p_{\boldsymbol{\theta}}(\boldsymbol{\theta})d\boldsymbol{\theta}} \tag{A.7}$$

where $p_{Y|\boldsymbol{\theta},U}(Y|\boldsymbol{\theta},U)$ is the likelihood function, that is, the probability of observing a certain sequence of CGM measurements $Y := \{y(t_k), t_k = k \cdot T_s, k = 1, \ldots, D\}$ given the parameter vector $\boldsymbol{\theta}$ and the input $U := \{\boldsymbol{u}(t_k), t_k = k \cdot T_s, k = 1, \ldots, D\}$ with D the number of available data points and T_s the sampling period, while $p_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ is the *a priori* information on the distributions of unknown parameters, which was obtained from previous studies [108][109][116].

The successful application of MCMC to address the *a priori* non-identifiability of the ReplayBG model was already discussed in [107].

In this work, an improved variant of the identification approach proposed in [107] was implemented, in which the Monte Carlo integration step is replaced by a new step where the *a posteriori* distribution $p_{\boldsymbol{\theta}|Y,U}(\boldsymbol{\theta}|Y,U)$ is fit and represented in sampled form.

From the practical point-of-view, $\boldsymbol{\theta}_i$ samples are generated from a Markov Chain whose stationary distribution is exactly the posterior in (A.7) (target distribution). Such a chain is built in MATLAB (Mathworks Inc., Natick, MA, USA) using an Adaptive Single Component Metropolis-Hastings scheme [117].

The obtained samples $\boldsymbol{\theta}_i$ are then used to fit a multivariate t copula distribution able to represent $p_{\boldsymbol{\theta}|Y,U}(\boldsymbol{\theta}|Y,U)$ and capture the underneath dependence between unknown parameters [118]. Finally, such a copula distribution is used to generate 1000 realizations of $\boldsymbol{\theta}$, i.e., $\hat{\boldsymbol{\theta}}_r, r = 1, \ldots, 1000$, representing in sampled form the target distribution $p_{\boldsymbol{\theta}|Y,U}(\boldsymbol{\theta}|Y,U)$. These samples are stored for

their later use in the second step of ReplayBG.

A.3 Use of the model for simulation

The second step of ReplayBG uses the model identified in Step 1 to predict the hypothetical glucose concentration profile (which corresponds, in the model, to the variable IG, i.e., a noise-free version of CGM) that would have been obtained, in the same individual and in the same time window, from the adoption of an alternative therapy. Specifically, for each r-th realization of the model parameters, $\hat{\boldsymbol{\theta}}_r$, a corresponding $IG(t)$ profile, $\hat{IG}_r(t)$, is obtained by simulating the model (A.6) using as input the carbohydrate intakes and insulin administrations provided by the alternative therapy to be evaluated. As results, a total of 1000 predicted IG profile realizations are obtained. These realizations are ultimately used to infer the median IG profile, $\hat{IG}(t)_{50}$, and the boundaries of the corresponding $25^{th} - 75^{th}$ percentile confidence interval, $\hat{IG}(t)_{25}$, $\hat{IG}(t)_{75}$, of the therapy under assessment.

Bibliography

[1] K. Alberti and P. Zimmet, "Definition, diagnosis and classification of diabetes mellitus and its complications. Part 1: Diagnosis and classification of diabetes mellitus. Provisional report of a WHO Consultation," *Diabetic Medicine*, vol. 15, no. 7, pp. 539–553, 7 1998.

[2] D. Daneman, "Type 1 diabetes," *The Lancet*, vol. 367, no. 9513, pp. 847–858, 2006.

[3] American Diabetes Association, "Diagnosis and Classification of Diabetes Mellitus," *Diabetes Care*, vol. 33, no. Suppl 1, pp. S62–S69, 2010.

[4] P. E. Cryer, S. N. Davis, and H. Shamoon, "Hypoglycemia in Diabetes," *Diabetes Care*, vol. 26, no. 6, pp. 1902–1912, 2003.

[5] D. M. Nathan, "The Diabetes Control and Complications Trial/Epidemiology of Diabetes Interventions and Complications Study at 30 Years: Overview," *Diabetes Care*, vol. 37, no. 1, pp. 9–16, 2014.

[6] A. Janež, C. Guja, A. Mitrakou, N. Lalic, T. Tankova, L. Czupryniak, A. G. Tabák, M. Prazny, E. Martinka, and L. Smircic-Duvnjak, "Insulin therapy in adults with type 1 diabetes mellitus: a narrative review," *Diabetes Therapy*, vol. 11, no. 2, pp. 387–409, 2020.

[7] V. Valla, "Continuous subcutaneous insulin infusion (csii) pumps," *Advances in Experimental Medicine and Biology*, vol. 771, pp. 414–419, 2012.

[8] J. Kesavadev, B. Saboo, M. B. Krishna, and G. Krishnan, "Evolution of insulin delivery devices: from syringes, pens, and pumps to diy artificial pancreas," *Diabetes Therapy*, vol. 11, no. 6, pp. 1251–1269, 2020.

[9] A. D. Association, "American Diabetes Association: standard of medical care in diabetes," *Diabetes Care*, vol. 42, no. 1, pp. S61–S70, 2019.

[10] S. Schmidt and K. Nørgaard, "Bolus Calculators," *Journal of Diabetes Science and Technology*, vol. 8, no. 5, pp. 1035–1041, 2014.

[11] J. Walsh and R. Roberts, *Pumping insulin: everything you need for success on a smart insulin pump*. Torrey Pines Press San Diego, CA, 2006, vol. 4.

[12] T. M. Gross, D. Kayne, A. King, C. Rother, and S. Juth, "A Bolus Calculator Is an Effective Means of Controlling Postprandial Glycemia in Patients on Insulin Pump Therapy," *Diabetes Technology & Therapeutics*, vol. 5, no. 3, pp. 365–369, 2003.

[13] K. K. Trout, C. Homko, and N. C. Tkacs, "Methods of measuring insulin sensitivity," *Biological research for nursing*, vol. 8, no. 4, pp. 305–318, 2007.

[14] R. W. Beck, R. M. Bergenstal, L. M. Laffel, and J. C. Pickup, "Advances in technology for management of type 1 diabetes," *The Lancet*, vol. 394, no. 10205, pp. 1265–1273, 2019.

[15] G. Cappon et al., "Continuous Glucose Monitoring Sensors for Diabetes Management: A Review of Technologies and Applications," *Diabetes & Metabolism*, vol. 43, p. 383–397, Aug 2019.

[16] D. De Salvo and B. Buckingham, "Continuous glucose monitoring: current use and future directions." *Current Diabetes Reports*, vol. 13, no. 5, pp. 657–62, Oct 2013.

[17] M. Vettoretti, G. Cappon, G. Acciaroli, A. Facchinetti, and G. Sparacino, "Continuous Glucose Monitoring: Current Use in Diabetes Management and Possible Future Applications," *Journal of Diabetes Science and Technology*, vol. 12, no. 5, pp. 1064–1071, 2018.

[18] M. Montagnana, M. Caputo, D. Giavarina, and G. Lippi, "Overview on self-monitoring of blood glucose," *Clinica Chimica Acta*, vol. 402, no. 1-2, pp. 7–13, 2009.

[19] S. F. Clarke and J. R. Foster, "A History of Blood Glucose Meters and Their Role in Self-Monitoring of Diabetes Mellitus," *British Journal of Biomedical Science*, vol. 69, no. 2, pp. 83–93, 2012.

[20] M. Christiansen, T. Bailey, E. Watkins, D. Liljenquist, D. Price, K. Nakamura, R. Boock, and T. Peyser, "A new-generation continuous glucose monitoring system: improved accuracy and reliability compared with a previous-generation system," *Diabetes technology & therapeutics*, vol. 15, no. 10, pp. 881–888, 2013.

[21] J. R. Castle and P. G. Jacobs, "Nonadjunctive use of continuous glucose monitoring for diabetes treatment decisions," *Journal of diabetes science and technology*, vol. 10, no. 5, pp. 1169–1173, 2016.

[22] S. Marden, P. Thomas, Z. Sheppard, J. Knott, J. Lueddeke, and D. Kerr, "Poor numeracy skills are associated with glycaemic control in type 1 diabetes," *Diabetic Medicine*, vol. 29, no. 5, pp. 662–669, 2012.

[23] J. Walsh, R. Roberts, and T. Bailey, "Guidelines for optimal bolus calculator settings in adults," *Journal of diabetes science and technology*, vol. 5, no. 1, pp. 129–135, 2011.

[24] A. B. King, A. Kuroda, M. Matsuhisa, and T. Hobbs, "A review of insulin-dosing formulas for continuous subcutaneous insulin infusion (csii) for adults with type 1 diabetes," *Current diabetes reports*, vol. 16, no. 9, pp. 1–8, 2016.

[25] R. Harbison, M. Hecht, and J. MacLeod, "Building a data-driven multiple daily insulin therapy model using smart insulin pens," *Journal of Diabetes Science and Technology*, vol. 16, no. 3, pp. 610–616, 2022.

[26] L. Heinemann, G. Freckmann, D. Ehrmann, G. Faber-Heinemann, S. Guerra, D. Waldenmaier, and N. Hermanns, "Real-time continuous glucose monitoring in adults with type 1 diabetes and impaired hypoglycaemia awareness or severe hypoglycaemia treated with multiple daily insulin injections (hypode): a multicentre, randomised controlled trial," *The Lancet*, vol. 391, no. 10128, pp. 1367–1377, 2018.

[27] L. M. Norlander, S. Anderson, C. J. Levy, L. Ekhlaspour, D. W. Lam, L. Hsu, S. E. Loebner, S. J. Ogyaadu, G. O'MALLEY, C. M. Levister *et al.*, "Late and missed meal boluses with multiple daily insulin injections," *Diabetes*, vol. 67, no. Supplement_1, 2018.

[28] D. T. Ahn, "Automated bolus calculators and connected insulin pens: A smart combination for multiple daily injection insulin therapy," *Journal of Diabetes Science and Technology*, vol. 16, no. 3, pp. 605–609, 2022.

[29] T. S. Bailey and J. Y. Stone, "A novel pen-based bluetooth-enabled insulin delivery system with insulin dose tracking and advice," *Expert Opinion on Drug Delivery*, vol. 14, no. 5, pp. 697–703, 2017.

[30] B. W. Gildon, "Inpen smart insulin pen system: product review and user experience," *Diabetes Spectrum*, vol. 31, no. 4, pp. 354–358, 2018.

[31] T. M. Gross, D. Kayne, A. King, C. Rother, and S. Juth, "A bolus calculator is an effective means of controlling postprandial glycemia in patients on insulin pump therapy," *Diabetes technology & therapeutics*, vol. 5, no. 3, pp. 365–369, 2003.

[32] S. K. Garg, T. R. Bookout, K. K. McFann, W. C. Kelly, C. Beatson, S. L. Ellis, R. S. Gutin, and P. A. Gottlieb, "Improved glycemic control in intensively treated adult subjects with type 1 diabetes using insulin guidance software," *Diabetes technology & therapeutics*, vol. 10, no. 5, pp. 369–375, 2008.

[33] R. Hovorka, V. Canonico, L. J. Chassin, U. Haueter, M. Massi-Benedetti, M. Orsini Federici, T. R. Pieber, H. C. Schaller, L. Schaupp, T. Vering, and M. E. Wilinska, "Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes," *Physiological Measurement*, vol. 25, no. 4, pp. 905–920, 2004.

[34] M. E. Wilinska, L. J. Chassin, C. L. Acerini, J. M. Allen, D. B. Dunger, and R. Hovorka, "Simulation environment to evaluate closed-loop insulin delivery systems in type 1 diabetes," *Journal of Diabetes Science and Technology*, vol. 4, no. 1, pp. 132–144, 2010.

[35] S. S. Kanderian, S. Weinzimer, G. Voskanyan, and G. M. Steil, "Identification of intraday metabolic profiles during closed-loop glucose control in individuals with type 1 diabetes," *Journal of Diabetes Science and Technology*, vol. 3, no. 5, pp. 1047–1057, 2009.

[36] A. Haidar, M. E. Wilinska, J. A. Graveston, and R. Hovorka, "Stochastic virtual population of subjects with type 1 diabetes for the assessment

of closed-loop glucose controllers," *IEEE Transactions on Biomedical Engineering*, vol. 60, no. 12, pp. 3524–3533, 2013.

[37] R. Visentin, E. Campos-Náñez, M. Schiavon, D. Lv, M. Vettoretti, M. Breton, B. P. Kovatchev, C. Dalla Man, and C. Cobelli, "The uva/padova type 1 diabetes simulator goes from single meal to single day," *Journal of diabetes science and technology*, vol. 12, no. 2, pp. 273–281, 2018.

[38] C. Dalla Man, F. Micheletto, D. Lv, M. Breton, B. Kovatchev, and C. Cobelli, "The UVA/Padova Type 1 Diabetes Simulator: New Features," *Journal of Diabetes Science and Technology*, vol. 8, no. 1, pp. 26–34, 2014.

[39] 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," 2009.

[40] J. A. DiMasi, H. G. Grabowski, and R. W. Hansen, "The cost of drug development," *New England Journal of Medicine*, vol. 372, no. 20, pp. 1972–1972, 2015.

[41] M. Viceconti, C. Cobelli, T. Haddad, A. Himes, B. Kovatchev, and M. Palmer, "In silico assessment of biomedical products: the conundrum of rare but not so rare events in two case studies," *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, vol. 231, no. 5, pp. 455–466, 2017.

[42] P. Herrero, P. Pesl, M. Reddy, N. Oliver, P. Georgiou, and C. Toumazou, "Advanced insulin bolus advisor based on run-to-run control and case-based reasoning," *IEEE journal of biomedical and health informatics*, vol. 19, no. 3, pp. 1087–1096, 2014.

[43] C. Fabris, B. Ozaslan, and M. D. Breton, "Continuous glucose monitors and activity trackers to inform insulin dosing in type 1 diabetes: the university of virginia contribution," *Sensors*, vol. 19, no. 24, p. 5386, 2019.

[44] G. Cappon, M. Vettoretti, F. Marturano, A. Facchinetti, and G. Sparacino, "A neural-network-based approach to personalize insulin bolus calculation using continuous glucose monitoring," *Journal of Diabetes Science and Technology*, vol. 12, no. 2, pp. 265–272, 2018.

[45] I. Contreras, J. Vehi *et al.*, "Artificial intelligence for diabetes management and decision support: literature review," *Journal of medical Internet research*, vol. 20, no. 5, p. e10775, 2018.

[46] G. Noaro, G. Cappon, G. Sparacino, F. Boscari, D. Bruttomesso, and A. Facchinetti, "Methods for insulin bolus adjustment based on the continuous glucose monitoring trend arrows in type 1 diabetes: Performance and safety assessment in an in silico clinical trial," *Journal of Diabetes Science and Technology*, 2021.

[47] D. Bruttomesso, F. Boscari, G. Lepore, G. Noaro, G. Cappon, A. Girelli, L. Bozzetto, A. Tumminia, G. Grassi, G. Sparacino *et al.*, "A "slide rule" to adjust insulin dose using trend arrows in adults with type 1 diabetes: test in silico and in real life," *Diabetes Therapy*, vol. 12, no. 5, pp. 1313–1324, 2021.

[48] G. Noaro, G. Cappon, M. Vettoretti, G. Sparacino, S. Del Favero, and A. Facchinetti, "Machine-learning based model to improve insulin bolus calculation in type 1 diabetes therapy," *IEEE Transactions on Biomedical Engineering*, vol. 68, no. 1, pp. 247–255, 2020.

[49] G. Noaro, G. Cappon, G. Sparacino, S. Del Favero, and A. Facchinetti, "Nonlinear machine learning models for insulin bolus estimation in type

1 diabetes therapy," in *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2020, pp. 5502–5505.

[50] G. Noaro, G. Cappon, G. Sparacino, and A. Facchinetti, "An ensemble learning algorithm based on dynamic voting for targeting the optimal insulin dosage in type 1 diabetes management," in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2021, pp. 1828–1831.

[51] J. Pettus and S. V. Edelman, "Use of glucose rate of change arrows to adjust insulin therapy among individuals with type 1 diabetes who use continuous glucose monitoring," *Diabetes Technology & Therapeutics*, vol. 18, no. S2, pp. S2–34, 2016.

[52] G. Cappon, F. Marturano, M. Vettoretti, A. Facchinetti, and G. Sparacino, "In silico assessment of literature insulin bolus calculation methods accounting for glucose rate of change," *Journal of diabetes science and technology*, vol. 13, no. 1, pp. 103–110, 2019.

[53] G. Scheiner, *Practical CGM: A Guide to Improving Outcomes Through Continuous Glucose Monitoring*. Am Diabetes Assoc, 2015.

[54] R. Ziegler, S. von Sengbusch, J. Kröger, O. Schubert, P. Werkmeister, D. Deiss, and T. Siegmund, "Therapy Adjustments Based on Trend Arrows Using Continuous Glucose Monitoring Systems," *Journal of Diabetes Science and Technology*, vol. 13, no. 4, pp. 763–773, 2019.

[55] G. Aleppo, L. M. Laffel, A. J. Ahmann, I. B. Hirsch, D. F. Kruger, A. Peters, R. S. Weinstock, and D. R. Harris, "A Practical Approach to Using Trend Arrows on the Dexcom G5 CGM System for the Management of

Adults With Diabetes," *Journal of the Endocrine Society*, vol. 1, no. 12, pp. 1445–1460, Dec 2017.

[56] J. Pettus and S. V. Edelman, "Recommendations for Using Real-Time Continuous Glucose Monitoring (rtCGM) Data for Insulin Adjustments in Type 1 Diabetes," *Journal of Diabetes Science and Technology*, vol. 11, no. 1, pp. 138–147, 2017.

[57] B. Buckingham, D. Xing, S. Weinzimer, R. Fiallo-Scharer, C. Kollman, N. Mauras, E. Tsalikian, W. Tamborlane, T. Wysocki, K. Ruedy, and R. Beck, "Use of the DirecNet Applied Treatment Algorithm (DATA) for Diabetes Management with a Real-Time Continuous Glucose Monitor (the FreeStyle Navigator)," *Pediatric Diabetes*, vol. 9, pp. 142–147, Apr 2008.

[58] D. C. Klonoff and D. Kerr, "A Simplified Approach Using Rate of Change Arrows to Adjust Insulin With Real-Time Continuous Glucose Monitoring," *Journal of Diabetes Science and Technology*, vol. 11, no. 6, pp. 1063–1069, 2017.

[59] J. Pettus, D. A. Price, and S. V. Edelman, "How patients with type 1 diabetes translate continuous glucose monitoring data into diabetes management decisions," *Endocrine Practice*, vol. 21, no. 6, pp. 613–620, 2015.

[60] C. D. Man, F. Micheletto, D. Lv, M. Breton, B. Kovatchev, and C. Cobelli, "The uva/padova type 1 diabetes simulator: new features," *Journal of diabetes science and technology*, vol. 8, no. 1, pp. 26–34, 2014.

[61] T. Zhu, K. Li, L. Kuang, P. Herrero, and P. Georgiou, "An insulin bolus advisor for type 1 diabetes using deep reinforcement learning," *Sensors*, vol. 20, no. 18, p. 5058, 2020.

[62] T. Zhu, K. Li, P. Herrero, and P. Georgiou, "Basal glucose control in type 1 diabetes using deep reinforcement learning: An in silico validation," *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 4, pp. 1223–1232, 2021.

[63] B. P. Kovatchev, D. J. Cox, L. A. Gonder-Frederick, and W. Clarke, "Symmetrization of the blood glucose measurement scale and its applications," *Diabetes Care*, vol. 20, no. 11, pp. 1655–1658, 1997.

[64] T. Danne, R. Nimri, T. Battelino, R. M. Bergenstal, K. L. Close, J. H. DeVries, S. Garg, L. Heinemann, I. Hirsch, S. A. Amiel, R. Beck, E. Bosi, B. Buckingham, C. Cobelli, E. Dassau, F. J. Doyle, S. Heller, R. Hovorka, W. Jia, T. Jones, O. Kordonouri, B. Kovatchev, A. Kowalski, L. Laffel, D. Maahs, H. R. Murphy, K. Nørgaard, C. G. Parkin, E. Renard, B. Saboo, M. Scharf, W. V. Tamborlane, S. A. Weinzimer, and M. Phillip, "International Consensus on Use of Continuous Glucose Monitoring," *Diabetes Care*, vol. 40, no. 12, pp. 1631–1640, 2017.

[65] D. M. Maahs, B. A. Buckingham, J. R. Castle, A. Cinar, E. R. Damiano, E. Dassau, J. H. DeVries, F. J. Doyle, S. C. Griffen, A. Haidar, L. Heinemann, R. Hovorka, T. W. Jones, C. Kollman, B. Kovatchev, B. L. Levy, R. Nimri, D. N. O'Neal, M. Philip, E. Renard, S. J. Russell, S. A. Weinzimer, H. Zisser, and J. W. Lum, "Outcome Measures for Artificial Pancreas Clinical Trials: A Consensus Report," *Diabetes Care*, vol. 39, no. 7, pp. 1175–1179, 2016.

[66] B. P. Kovatchev, D. J. Cox, L. A. Gonder-Frederick, and W. Clarke, "Symmetrization of the blood glucose measurement scale and its applications," *Diabetes care*, vol. 20, no. 11, pp. 1655–1658, 1997.

[67] W. Clarke and B. Kovatchev, "Statistical Tools to Analyze Continuous Glucose Monitor Data," *Diabetes Technology & Therapeutics*, vol. 11 Suppl 1, pp. S45–54, Jun 2009.

[68] C. F. Dormann, J. Elith, S. Bacher, C. Buchmann, G. Carl, G. Carré, T. Diekötter, J. García Márquez, B. Gruber, B. Lafourcade, P. Leitão, T. Münkemüller, C. Mcclean, P. Osborne, B. Reineking, B. Schröder, A. Skidmore, D. Zurell, and S. Lautenbach, "Collinearity: A review of methods to deal with it and a simulation study evaluating their performance," *Ecography*, vol. 36, pp. 27–46, 04 2013.

[69] T. Hastie, R. Tibshirani, and J. H. Friedman, *The elements of statistical learning: data mining, inference, and prediction, 2nd Edition*, ser. Springer series in statistics. Springer, 2009.

[70] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.

[71] T. K. Ho, "The random subspace method for constructing decision forests," *IEEE transactions on pattern analysis and machine intelligence*, vol. 20, no. 8, pp. 832–844, 1998.

[72] A. Natekin and A. Knoll, "Gradient boosting machines, a tutorial," *Frontiers in neurorobotics*, vol. 7, p. 21, 2013.

[73] S. Puuronen, V. Terziyan, and A. Tsymbal, "A dynamic integration algorithm for an ensemble of classifiers," in *International symposium on methodologies for intelligent systems*. Springer, 1999, pp. 592–600.

[74] L. E. Peterson, "K-nearest neighbor," *Scholarpedia*, vol. 4, no. 2, p. 1883, 2009.

[75] J. Mendes-Moreira, A. M. Jorge, C. Soares, and J. F. d. Sousa, "Ensemble learning: A study on different variants of the dynamic selection approach," in *International Workshop on Machine Learning and Data Mining in Pattern Recognition*. Springer, 2009, pp. 191–205.

[76] M. Robnik-Šikonja and I. Kononenko, "Theoretical and empirical analysis of relieff and rrelieff," *Machine learning*, vol. 53, no. 1, pp. 23–69, 2003.

[77] G. Cappon, M. Vettoretti, G. Sparacino, S. Del Favero, and A. Facchinetti, "ReplayBG: a methodology to identify a personalized model from type 1 diabetes data and simulate glucose concentrations to assess alternative therapies," *IEEE Transactions on Biomedical Engineering*, under revision.

[78] J. Kropff, S. Del Favero, J. Place, C. Toffanin, R. Visentin, M. Monaro, M. Messori, F. Di Palma, G. Lanzola, A. Farret, F. Boscari, S. Galasso, P. Magni, A. Avogaro, P. Keith-Hynes, B. Kovatchev, D. Bruttomesso, C. Cobelli, J. DeVries, and L. Magni, "2 month evening and night closed-loop glucose control in patients with type 1 diabetes under free-living conditions: A randomised crossover trial," *Lancet Diabetes Endocrinol.*, vol. 3, 10 2015.

[79] R. N. Bergman, Y. Z. Ider, C. R. Bowden, and C. Cobelli, "Quantitative estimation of insulin sensitivity." *American Journal of Physiology-Endocrinology and Metabolism*, vol. 236, no. 6, pp. E667–E677, 1979.

[80] M. Schiavon, C. Dalla Man, and C. Cobelli, "Modeling subcutaneous absorption of fast-acting insulin in type 1 diabetes," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 9, pp. 2079–2086, Sep. 2018.

[81] C. Dalla Man, M. Camilleri, and C. Cobelli, "A system model of oral glucose absorption: Validation on gold standard data," *IEEE Transactions on Biomedical Engineering*, vol. 53, no. 12, pp. 2472–2478, Dec 2006.

[82] C. J. Watkins and P. Dayan, "Q-learning," *Machine learning*, vol. 8, no. 3, pp. 279–292, 1992.

[83] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[84] K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, "Deep reinforcement learning: A brief survey," *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 26–38, 2017.

[85] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double q-learning," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 30, no. 1, 2016.

[86] A. Facchinetti, S. Del Favero, G. Sparacino, and C. Cobelli, "Model of glucose sensor error components: identification and assessment for new dexcom g4 generation devices," *Medical & biological engineering & computing*, vol. 53, no. 12, pp. 1259–1269, 2015.

[87] M. Vettoretti, C. Battocchio, G. Sparacino, and A. Facchinetti, "Development of an error model for a factory-calibrated continuous glucose monitoring sensor with 10-day lifetime," *Sensors*, vol. 19, no. 23, p. 5320, 2019.

[88] M. Vettoretti, A. Facchinetti, G. Sparacino, and C. Cobelli, "Type-1 diabetes patient decision simulator for in silico testing safety and effectiveness of insulin treatments," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 6, pp. 1281–1290, 2017.

[89] A. Brazeau, H. Mircescu, K. Desjardins, C. Leroux, I. Strychar, J. Ekoé, and R. Rabasa-Lhoret, "Carbohydrate counting accuracy and blood glu-

cose variability in adults with type 1 diabetes," *Diabetes research and clinical practice*, vol. 99, no. 1, pp. 19–23, 2013.

[90] C. Roversi, M. Vettoretti, S. Del Favero, A. Facchinetti, G. Sparacino, and H.-R. Consortium, "Modeling carbohydrate counting error in type 1 diabetes management," *Diabetes technology & therapeutics*, vol. 22, no. 10, pp. 749–759, 2020.

[91] L. Kaufman and P. J. Rousseeuw, *Finding groups in data: an introduction to cluster analysis*. John Wiley & Sons, 2009.

[92] D. M. Maahs, B. A. Buckingham, J. R. Castle, A. Cinar, E. R. Damiano, E. Dassau, J. H. DeVries, F. J. Doyle III, S. C. Griffen, A. Haidar *et al.*, "Outcome measures for artificial pancreas clinical trials: a consensus report," *Diabetes care*, vol. 39, no. 7, pp. 1175–1179, 2016.

[93] T. Danne, R. Nimri, T. Battelino, R. M. Bergenstal, K. L. Close, J. H. DeVries, S. Garg, L. Heinemann, I. Hirsch, S. A. Amiel *et al.*, "International consensus on use of continuous glucose monitoring," *Diabetes care*, vol. 40, no. 12, pp. 1631–1640, 2017.

[94] W. J. Conover, *Practical nonparametric statistics*. john wiley & sons, 1999, vol. 350.

[95] R. Visentin, E. Campos-Náñez, M. Schiavon, D. Lv, M. Vettoretti, M. Breton, B. P. Kovatchev, C. D. Man, and C. Cobelli, "The UVA/Padova Type 1 Diabetes Simulator Goes From Single Meal to Single Day," *Journal of Diabetes Science and Technology*, vol. 12, no. 2, pp. 273–281, 2018.

[96] M. Vettoretti, A. Facchinetti, G. Sparacino, and C. Cobelli, "Type-1 Diabetes Patient Decision Simulator for In Silico Testing Safety and Effec-

tiveness of Insulin Treatments," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 6, pp. 1281–1290, 2018.

[97] A. Facchinetti, S. Del Favero, G. Sparacino, and C. Cobelli, "Model of glucose sensor error components: identification and assessment for new Dexcom G4 generation devices," *Medical & Biological Engineering & Computing*, vol. 53, pp. 1259–1269, Nov 2014.

[98] M. Vettoretti, C. Battocchio, G. Sparacino, and A. Facchinetti, "Development of an Error Model for a Factory-Calibrated Continuous Glucose Monitoring Sensor with 10-Day Lifetime," *Sensors*, vol. 19, p. 5320, Dec 2019.

[99] D. C. Klonoff, "The current status of mhealth for diabetes: will it be the next big thing?" *Journal of diabetes science and technology*, vol. 7, no. 3, pp. 749–758, 2013.

[100] P. J. O'Connor, J. M. Sperl-Hillen, C. J. Fazio, B. M. Averbeck, B. H. Rank, and K. L. Margolis, "Outpatient diabetes clinical decision support: current status and future directions," *Diabetic Medicine*, vol. 33, no. 6, pp. 734–741, 2016.

[101] P. Jia, P. Zhao, J. Chen, and M. Zhang, "Evaluation of clinical decision support systems for diabetes care: An overview of current evidence," *Journal of Evaluation in Clinical Practice*, vol. 25, no. 1, pp. 66–77, 2018.

[102] R. Nimri, A. R. Ochs, J. E. Pinsker, M. Phillip, and E. Dassau, "Decision support systems and closed loop," *Diabetes Technology & Therapeutics*, vol. 21, no. S1, pp. S–42, 2019.

[103] C. Marling, J. Shubrook, and F. Schwartz, "Toward case-based reasoning for diabetes management: A preliminary clinical study and decision

support system prototype," *Computational intelligence*, vol. 25, no. 3, pp. 165–179, 2009.

[104] N. S. Tyler, C. M. Mosquera-Lopez, L. M. Wilson, R. H. Dodier, D. L. Branigan, V. B. Gabo, F. H. Guillot, W. W. Hilts, J. El Youssef, J. R. Castle *et al.*, "An artificial intelligence decision support system for the management of type 1 diabetes," *Nature metabolism*, vol. 2, no. 7, pp. 612–619, 2020.

[105] N. S. Tyler and P. G. Jacobs, "Artificial intelligence in decision support systems for type 1 diabetes," *Sensors*, vol. 20, no. 11, p. 3214, 2020.

[106] G. Cappon, G. Noaro, N. Camerlingo, L. Cossu, G. Sparacino, and A. Facchinetti, "A new decision support system for type 1 diabetes management," in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2021, pp. 1993–1996.

[107] G. Cappon, A. Facchinetti, G. Sparacino, and S. Del Favero, "A Bayesian Framework to Identify Type 1 Diabetes Physiological Models Using Easily Accessible Patient Data," in *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Jul 2019, pp. 6914–6917.

[108] M. Schiavon, C. Dalla Man, and C. Cobelli, "Modeling subcutaneous absorption of fast-acting insulin in type 1 diabetes," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 9, pp. 2079–2086, 2018.

[109] C. Dalla Man, M. Camilleri, and C. Cobelli, "A system model of oral glucose absorption: Validation on gold standard data," *IEEE Transactions on Biomedical Engineering*, vol. 53, no. 12, pp. 2472–2478, 2006.

[110] R. N. Bergman, C. R. Bowden, and C. Cobelli, "Quantitative estimation of insulin sensitivity," *American Physiological Society Journal*, vol. 236, no. 6, pp. e667–e677, 1979.

[111] C. Dalla Man, A. Caoumo, and C. Cobelli, "The oral glucose minimal model: estimation of insulin sensitivity from a meal test," *IEEE Trans. Biomed. Eng.*, vol. 184, no. 1, pp. 419–429, 2002.

[112] C. Dalla Man, F. Micheletto, D. Lv, M. Breton, B. Kovatchev, and C. Cobelli, "The uva/padova type 1 diabetes simulator: New features," *Journal of Diabetes Science and Technology*, vol. 8, no. 1, pp. 26–34, 2017.

[113] A. Facchinetti, G. Sparacino, and C. Cobelli, "Reconstruction of glucose in plasma from interstitial fluid continuous glucose monitoring data: role of sensor calibration," *Journal of Diabetes Science and Technology*, vol. 1, no. 5, pp. 617–623, 2007.

[114] K. Rebrin, G. M. Steil, W. van Antwerp, and J. J. Mastrototaro, "Subcutaneous glucose predicts plasma glucose independent of insulin: implications for continuous monitoring," *American Physiological Society Journal*, vol. 277, no. 3, pp. E561–E571, 1999.

[115] G. Pillonetto, G. Sparacino, and C. Cobelli, "Numerical non-identifiability regions of the minimal model of glucose kinetics: superiority of bayesian estimation," *Mathematical Biosciences*, vol. 184, no. 1, pp. 53–67, 2003.

[116] C. Dalla Man, A. Caumo, R. Basu, R. Rizza, G. Toffolo, and C. Cobelli, "Minimal model estimation of glucose absorption and insulin sensitivity from oral test: validation with a tracer method," *Journal of Theoretical Medicine*, vol. 1, no. 4, pp. 313–323, 1999.

[117] W. R. Gilks, S. Richardson, and D. Spiegelhalter, "Markov chain monte carlo in practice," *1st edition: Chapman and Hall*, 1996.

[118] S. Demarta and A. McNeil, "The t copula and related copulas," *International Statistical Review*, vol. 73, no. 1, pp. 111–129, 2007.