Available online at www.sciencedirect.com

ScienceDirect

IFAC Papers Online CONFERENCE PAPER ARCHIVE

IFAC PapersOnLine 56-2 (2023) 9642–9647

Artificial Pancreas under a Zone Model Predictive Control based on Gaussian Process models: toward the personalization of the closed loop

Marco Polver * **Beatrice Sonzogni** * **Mirko Mazzoleni** *
Fabio Previdi * **Antonio Ferramosca** *

* *Dipartimento di Ingegneria Gestionale, dell'Informazione e della Produzione, Università degli Studi di Bergamo (Italy)*

Abstract: This work introduces a novel zone model predictive control (MPC) based on Gaussian Process models (GPs) for the artificial pancreas (AP). The main novelty of the proposal is to exploit a GP that is trained on previously collected metabolic data of type 1 diabetes mellitus (T1DM) patients, to regulate the blood glucose levels by means of a personalized MPC strategy that automatically adjusts the basal insulin and the insulin boluses to be injected to the patients. The average closed-loop performance is improved in terms of classical indexes such as time in range, avoidance of critic hypoglycaemia episodes and avoidance of long-term hyperglycaemia events. The controller was evaluated in-silico by means of the FDA-accepted UVA/Padova metabolic simulator on 11 adult T1DM patients, showing promising results.

Keywords: Artificial Pancreas, Model Predictive Control, Data-driven Control, Gaussian Processes

1. INTRODUCTION

Type 1 Diabetes Mellitus (T1DM) is a disease characterized by the progressive destruction of the beta cells in the islets of Langerhans of the pancreas, which are responsible for the production of insulin. The lack of insulin caused by this phenomenon results in the increase of glucose in the bloodstream.

In non-diabetic people, a low amount of insulin is continuously released to regulate the blood sugar overnight and between meals (basal insulin), while larger amounts of insulin (insulin boluses) are released after meals. The objective of every diabetes therapy is to mimic this behavior. A modern approach for treating T1DM is the usage of a commercial device, called artificial pancreas (AP), that regulates the glucose level by continuously measuring it with a sensor and using a controller to compute the right amount of insulin to be injected by means of an infusion pump (Moon et al., 2021).

Several controllers have been tested on real and in-silico patients, starting from more simple solutions such as PID controllers, like those proposed in (Huyett et al., 2015) and (Rosales et al., 2022), and then moving on to more complex strategies such as Model Predictive Control (MPC), see (Soru et al., 2012; Gondhalekar et al., 2016; González et al., 2017, 2020; Del Favero et al., 2019; Kovatchev, 2018; Hovorka et al., 2004; Boiroux et al., 2018; Hajizadeh et al., 2019; Shi et al., 2018; Toffanin et al., 2013).

In order to use MPC inside an AP, a model of the endocrine-metabolic system is needed. Though first-principle models have been developed for this purpose, their complexity is leading control system researchers to replace them with black-box models. Particular interest has been posed to the development of MPCs based on neural network models, like in (Bahremand et al., 2019) and (Dutta et al., 2018). Among the black-box models that have been used inside MPCs in other fields, we find also Gaussian Process models (GPs). This is because Gaussian Process Regression (GPR) provides prediction probability distributions whose variance can be exploited to make sure that state and output constraints are satisfied. However, in the case of APs, GPs have only been used inside sliding mode controllers (Patra and Rout, 2017) and reinforcement learning algorithms (De Paula et al., 2015) or to detect changes in the insulin sensitivity (IS) of a person (Ortmann et al., 2017).

In this article we propose a zone MPC for the control of the glucose level in T1DM patients which uses a GP identified on the individual patient as a personalized model of the endocrine-metabolic system. In particular, the GP is identified on historical data about the subject's metabolism and allows the MPC to automatically compute the right amount of basal insulin to be continuously provided and also the insulin boluses that have to be injected at mealtimes. The sampling time of the system is considered to be 10 minutes and the meals are assumed to be announced exactly at mealtime. The proposed controller was tested on the 11 virtual adult patients of the UVA/Padova simulator (The Epsilon Group, 2016), providing satisfactory results.

The rest of the article is organized as follows: Section 2 formulates the control problem and provides some preliminary concepts about GPs, Section 3 describes the proposed MPC and Section 4 presents the results obtained on in-silico patients. Conclusions drawn are given in Section 5.

⋆ This work was funded by the **National Plan for NRRP Complementary Investments** in the **call for the funding of research initiatives for technologies and innovative trajectories in the health and care sectors** - project **AdvaNced Technologies for Human-centrEd Medicine (ANTHEM)** CUP: B53C22006700001.

10.1016/j.ifacol.2023.10.271

Notation: Concatenations of column vectors are represented as $(a, b) := [a^\top \ b^\top]^\top$. For time-independent quantities, integer subscripts are used to distinguish between different measures, e.g. w_1. In the case of time-dependent quantities, realized quantities are time-indexed using parentheses, e.g. $y(k)$, while predicted quantities use subscripts, e.g. $y_{k+i|k}$ is an i-step-ahead prediction computed at time step k.

2. PRELIMINARIES

In this section, the problem to be solved is formulated and basic concepts about the usage of GPs and GPR for the identification of dynamical systems are given.

2.1 Problem formulation

The endocrine-metabolic system of T1DM patients, treated with an AP providing basal insulin and insulin boluses, can be seen as a nonlinear discrete-time multiple-input-single-output system. The output of this system is the blood glucose level measured by a continuous glucose monitor (CGM) and is considered to behave as a NARX model of the form

$$\begin{aligned} y(k) = f(&y(k-1), \dots, y(k-n_a), \\ &u_{\text{meal}}(k-1), \dots, u_{\text{meal}}(k-n_b), \\ &u_{\text{bolus}}(k-1), \dots, u_{\text{bolus}}(k-n_c), \\ &u_{\text{basal}}(k-1), \dots, u_{\text{basal}}(k-n_d)) + \epsilon(k), \end{aligned} \tag{1}$$

where k denotes the discrete-time index, $y(k) \in \mathbb{R}$ is corrupted by i.i.d. white Gaussian noise $\epsilon(k) \sim \mathcal{N}(0, \sigma_\epsilon^2)$, and with n_a, n_b, n_c and n_d being the memory horizons of glucose, meals, insulin boluses and basal insulin that are considered to build the NARX state.

While u_{bolus} and u_{basal} are controlled inputs, u_{meal} is an uncontrolled input. Moreover, meals are considered to be announced in a timely manner and without any time advance. The controlled inputs are subject to constraints that depend on the infusion pump's capabilities and the subject's insulin sensitivity:

$$(u_{\text{bolus}}(k), u_{\text{basal}}(k)) \in \mathbb{U}, \quad \forall k \geq 0. \tag{2}$$

The control objective is to keep the blood glucose level y inside the euglycaemic range, namely:

$$70\text{mg/dl} \leq y(k) \leq 180\text{mg/dl}, \quad \forall k \geq 0. \tag{3}$$

Historical data about the considered T1DM patient's blood glucose, meals, insulin boluses and basal insulin are available. These data are exploited to determine the hyperparameters of a GP that is used inside a zone MPC that has the goal to keep the blood glucose level within the boundaries defined in (3).

2.2 Gaussian Process models for dynamical systems

A GP is a collection of random variables, any finite number of which have consistent joint Gaussian distributions, and is completely determined by a mean function ϕ and a covariance function c, see (Rasmussen, 2003). For every finite number of input vectors $[w_1, w_2, \dots, w_n]^\top$, with $w_i \in \mathbb{R}^{n_w}$, the function values are assumed to behave according to:

$$\mathbf{f} = \begin{bmatrix} f(w_1) \\ \vdots \\ f(w_n) \end{bmatrix} \sim \mathcal{N}(\phi_f, C), \ \phi_f = \begin{bmatrix} \phi(w_1) \\ \vdots \\ \phi(w_n) \end{bmatrix},$$

$$C = \begin{bmatrix} c(w_1, w_1) & \dots & c(w_n, w_1) \\ \vdots & \ddots & \vdots \\ c(w_1, w_n) & \dots & c(w_n, w_n) \end{bmatrix}$$

where $\phi : \mathbb{R}^{n_w} \to \mathbb{R}$ is a mean function and $c : \mathbb{R}^{n_w} \times \mathbb{R}^{n_w} \to \mathbb{R}$ is a covariance function. Such a GP is indicated as $f \sim \mathcal{GP}(\phi, c)$. The choice of ϕ and c depends on the characteristics of the function that has to be modelled.

Every covariance function depends on a set of hyperparameters ζ, which have to be determined. Assuming to have noisy observations of the function f and the corresponding input vectors, namely $y_i = f(w_i) + \epsilon_i$, $i = 1, \dots, n$, where $\epsilon_i \sim \mathcal{N}(0, \sigma_\epsilon^2)$, the hyperparameters $\eta = (\zeta, \sigma_\epsilon^2)$ can be determined in different ways. In this paper, our choice is to determine them via maximization of the marginal likelihood $p(\mathbf{y}|W, \eta)$, where $\mathbf{y} = (y_1, \dots, y_n)$.

GPR assumes that $f \sim \mathcal{GP}(\phi, c)$ and uses a Bayesian approach to compute the predictive distribution of the noisy output y_* at a test input vector w_* given a training set composed of n pairs (w_i, y_i):

$$p(y_*|\mathbf{y}, w_*) = \mathcal{N}(\mu(w_*), \sigma^2(w_*))$$

$$\mu(w_*) = \phi(w_*) + c_*^\top (C + \sigma_\epsilon^2 I_n)^{-1}(\mathbf{y} - \phi_f) = \phi(w_*) + c_*^\top \chi \tag{4}$$

$$\begin{aligned} \sigma^2(w_*) &= c(w_*, w_*) - c_*^\top (C + \sigma_\epsilon^2 I_n)^{-1} c_* + \sigma_\epsilon^2 \\ &= c(w_*, w_*) - c_*^\top \psi + \sigma_\epsilon^2 \end{aligned} \tag{5}$$

where $\mathbf{y} = (y_1, \dots, y_n)$, $c_* = (c(w_1, w_*), \dots, c(w_n, w_*))$, $C \in \mathbb{R}^{n \times n}$ is a covariance matrix such that $[C]_{ij} = c(w_j, w_i)$ $\forall i, j = 1, \dots, n$, $\phi_f = (\phi(w_1), \dots, \phi(w_n))$, $\chi = (C + \sigma_\epsilon^2 I_n)^{-1}(\mathbf{y} - \phi_f)$ and $\psi = (C + \sigma_\epsilon^2 I_n)^{-1} c_*$.

GPs can be easily adapted to identify dynamical systems by building regressors that are made of past outputs and inputs. In this way, GPR can be used to compute one-step-ahead predictions of the system's output. Within an MPC, however, multi-step-ahead predictions are needed, leading to the issue of uncertainty propagation. This comes from the fact that standard GPR works on deterministic regressors, while multi-step-ahead predictions require to use GPR on stochastic regressors, as they also include past predictions, which however are Gaussian.

Given that ignoring the stochasticity of such regressors would lead to predictions with underestimated variances, in this paper we assume that the posterior predictive distribution computed over stochastic regressors is still Gaussian and compute its mean and variance by means of the law of iterated expectations and the law of conditional variances explained in (Quinonero-Candela et al., 2003).

3. CONTROLLER DESIGN

In this section, the proposed MPC is described in detail and the entire design process is presented.

3.1 A Gaussian Process model of the endocrine-metabolic system

As stated in the previous section, historical data about the considered T1DM patient are assumed to be available. In particular, the available dataset takes the form:

$$\mathcal{D} := \{(y(i), u_{\text{meal}}(i), u_{\text{bolus}}(i), u_{\text{basal}}(i), \quad i = 1, \dots, n_{\mathcal{D}})\}.$$

These data have to be used to determine the hyperparameters of the GP that will model the endocrine-metabolic system inside the MPC.

However, before determining such hyperparameters, it is necessary to choose the mean function ϕ and the covariance function c that characterize the GP and the memory horizons n_a, n_b, n_c and n_d. Moreover, since GPR requires inverting the matrix $C + \sigma_\epsilon^2 I_n \in \mathbb{R}^{n \times n}$, it can become

computationally demanding if n is large. For this reason, it is also convenient to train the GP on a subset of the initially available dataset, in order to have $n \ll n_{\mathcal{D}}$.
In this work, we choose to set $\phi(w) = 0$ and to use the squared exponential covariance function with automatic relevance determination (ARD), namely:

$$c(w_1, w_2|\zeta) = e^{-\frac{1}{2}(w_1-w_2)^\top \Lambda^{-1}(w_1-w_2)}, \tag{6}$$

where $\Lambda = \mathrm{diag}(\lambda_1^2, \ldots, \lambda_{n_d}^2)$ and $\zeta = (\lambda_1, \ldots, \lambda_{n_d})$ are the covariance hyperparameters Rasmussen (2003). The choice of this covariance function is motivated by the fact that the ARD function induces a RKHS space such that the estimated functions are BIBO stable nonlinear dynamical systems (Pillonetto, 2018).
The choice of the memory horizons n_a, n_b, n_c and n_d and of the number n of observations to keep requires further attention. Since the dynamics of the endocrine-metabolic system can differ significantly from patient to patient, these parameters have to be chosen for each person.
While a straightforward way of finding the optimal value of these parameters might be to train different models on a training set and choose as the final model the one providing the maximum marginal likelihood over a validation set, this method might bring unsatisfactory results. This is due to the fact that a GP can fit the data well even if it misunderstands the effect that each regressor's component has on the output function. In particular, since insulin boluses are often provided at mealtimes, many models providing good marginal likelihood values actually interpret boluses as inputs that cause an increase in blood glucose and meals as inputs causing a decrease.
For this reason, provided that the parameters to be found belong to the set of natural numbers and can take only a finite number of values, a better sequence of actions that can be accomplished for model selection is the following:

(1) For each combination of n_a, n_b, n_c, n_d and n:
   (a) The original dataset is transformed into:

$$\begin{aligned}\mathcal{D}_{\text{temp}} := \{&(y(i), w(i)),\\ &i = \max\{n_a, n_b, n_c, n_d\} + 1, \ldots, n_{\mathcal{D}}\},\end{aligned}$$

   where

$$\begin{aligned}w(i) := (&y(i-1), \ldots, y(i-n_a),\\ &u_{\text{meal}}(i-1), \ldots, u_{\text{meal}}(i-n_b),\\ &u_{\text{bolus}}(i-1), \ldots, u_{\text{bolus}}(i-n_c),\\ &u_{\text{basal}}(i-1), \ldots, u_{\text{basal}}(i-n_d)).\end{aligned} \tag{7}$$

   (b) A subset of $\mathcal{D}_{\text{temp}}$ is built by selecting the n most informative observations. In this work this selection is done by running a k-means clustering algorithm with k set to n and then keeping the observations that are closest to the clusters' centroids.
   (c) The hyperparameters of the GP are found by maximization of the log marginal likelihood.
(2) The 10 models providing the best marginal likelihood values over a validation set are manually evaluated in order to choose the most consistent with the known effects of the system's inputs. In particular, a good model is expected to predict a decrease in the blood glucose level when the value of the provided basal insulin is large and to predict an increase after meals, which should be smaller when also an insulin bolus is provided at mealtime.

3.2 The optimization problem

The goal of the here presented controller is to keep the blood glucose level of T1DM patients within the euglycaemic range (3), while trying to minimize the amount of injected insulin. At each sampling instant k, the controller has to compute the best $u_{\text{basal}}(k)$ and $u_{\text{bolus}}(k)$ based on $y(k)$, a collection of past input and output values $w(k)$, where $w(k)$ has the same structure introduced in (7), and $u_{\text{meal}}(k)$.
The predictions of the blood glucose levels along the prediction horizon are carried out via GPR. At each prediction step j, a regressor is built such that

$$\begin{aligned}w_{k+j|k} := (&y_{k+j-1|k}, \ldots, y_{k+j-n_a|k},\\ &u_{\text{meal},k+j-1|k}, \ldots, u_{\text{meal},k+j-n_b|k},\\ &u_{\text{bolus},k+j-1|k}, \ldots, u_{\text{bolus},k+j-n_c|k},\\ &u_{\text{basal},k+j-1|k}, \ldots, u_{\text{basal},k+j-n_d|k}),\end{aligned}$$

where:

- $y_{k+j-i|k} = y(k+j-i)$ for $j \le i$;
- $y_{k+j-i|k} \sim \mathcal{N}(\mathsf{y}_{k+j-i|k}, \sigma^2_{k+j-i|k})$ for $j > i$, as they are the result of previous predictions.

The regressor $w_{k+j|k}$ is then used to get the prediction $y_{k+j|k}$, whose mean and variance are computed as explained in (Quinonero-Candela et al., 2003).
Given that hyperglycaemia and hypoglycaemia events, although undesirable, are possible and that the latter are more critical than the former, the cost function and the constraints involved in our MPC are inspired by (Abuin et al., 2020). In particular, soft constraints are provided to keep the predicted blood glucose levels within the euglycaemic range if possible, and the related slack variables are weighted inside the cost function in such a way that makes hypoglycaemia events less likely than hyperglycaemia ones, because they are more dangerous.
The soft constraints on the predicted blood glucose levels are made even safer by considering also the standard deviation of each prediction. This comes from the fact that both predictions with expected values that lie outside the euglycaemic range and predictions with large variances are to be considered unsafe.
Moreover, our MPC belongs to the family of zone MPCs, as it brings the system's output to a specified range, or zone, instead of bringing it to a reference. This is done by introducing a new optimization variable y_a, which is imposed to stay within a desired range, and by weighting the distance between the predicted blood glucose levels and this variable in the cost function.
Along the chosen prediction horizon N, basal insulin is free to change at every prediction step. The same doesn't apply to insulin boluses; a bolus can be provided only at the first prediction step and only when $u_{\text{meal}} > 0$.
The last concepts and assumptions formally translate into the following optimization problem, which has to be solved at each sampling instant k:

$$\begin{aligned}\min_{\mathbf{u}} \sum_{j=0}^{N} &\Big[Q(\mathsf{y}_{k+j|k} - y_a)^2 + \rho_{\text{hypo}}\delta^2_{\text{hypo},k+j} + \rho_{\text{hyper}}\delta^2_{\text{hyper},k+j}\Big]\\ &+ \sum_{j=0}^{N-1} \Big[R_{\text{basal}} u^2_{\text{basal},k+j}\Big] + R_{\text{bolus}} u^2_{\text{bolus},k}\end{aligned} \tag{8a}$$

$$\text{s.t. } y_{k|k} = y(k) \tag{8b}$$

$$y_{k+j|k} \sim \mathcal{N}(\mathsf{y}_{k+i|k}, \sigma^2_{k+i|k}) \quad \forall j \in [1, N] \tag{8c}$$
$$0 \leq u_{\text{basal},k+j} \leq \overline{u}_{\text{basal}} \quad \forall j \in [0, N-1] \tag{8d}$$
$$0 \leq u_{\text{bolus},k} \leq \overline{u}_{\text{bolus}}(u_{\text{meal}}(k)) \tag{8e}$$
$$\underline{y} \leq y_a \leq (\underline{y} + \overline{y})/2 \tag{8f}$$
$$\delta_{\text{hypo},k+j} \geq 0,\ \delta_{\text{hyper},k+j} \geq 0\ \forall j \in [0, N] \tag{8g}$$
$$\mathsf{y}_{k+j|k} - n_\sigma \sigma_{k+j|k} \geq \underline{y} - \delta_{\text{hypo},k+j}\ \forall j \in [0, N] \tag{8h}$$
$$\mathsf{y}_{k+j|k} + n_\sigma \sigma_{k+j|k} \leq \overline{y} + \delta_{\text{hyper},k+j}\ \forall j \in [0, N] \tag{8i}$$

where:

- $y_{k+j|k}$ are Gaussian distributions obtained by means of multi-step ahead prediction.
- $\overline{u}_{\text{basal}}$ is a design parameter and depends on the insulin sensitivity of the subject.
- The insulin boluses upper bound is computed as
$$\overline{u}_{\text{bolus}}(u_{\text{meal}}(k)) = \begin{cases} 0, & u_{\text{meal}}(k)) = 0 \\ \overline{\overline{u}}_{\text{bolus}}, & u_{\text{meal}}(k)) > 0 \end{cases}$$
where $\overline{\overline{u}}_{\text{bolus}}$ is a design parameter.
- $\delta_{\text{hypo},k+j}$ and $\delta_{\text{hyper},k+j}$ are the slack variables related to hypoglycaemia and hyperglycaemia events.
- $\underline{y} = 70$mg/dl and $\overline{y} = 180$mg/dl are the lower and upper bounds of the euglycaemic range. Note that in this work, y_a is imposed to stay within the lower half of the euglycaemic range $[70, 125]$mg/dl instead of the entire euglycaemic range $[70, 180]$mg/dl to recover more quickly from possible hyperglycaemia events.
- n_σ is a design parameter. The higher n_σ, the more our MPC avoids uncertain predictions.

4. IN-SILICO EVALUATION

The designed zone MPC was tested on 11 in-silico adult patients by means of the UVA/Padova T1DM simulator implemented in the software DMMS.R. The simulations were run on a machine equipped with an Intel® Xeon® E3-1225 CPU and 8GB of RAM. The controller was implemented in MATLAB® and the optimization problem that is inherent in each MPC step was solved by means of the `fmincon` function.

Both the data collection phase and the final simulations were carried out considering a sampling time $T_s = 10$min. This choice is motivated by the need to guarantee that the execution time of a single MPC step is shorter than the sampling time. In order to create GPs that really understand the dynamics of the endocrine-metabolic system, it is necessary to create regressors that take into account output and input values of the past 1-2 hours; for this reason choosing a shorter sampling time would lead to the need of building regressors with larger memory horizons n_a, n_b, n_c and n_d, thus making the optimization problem involved in the here proposed MPC computationally demanding.

In both the data collection phase and the final simulations, the daily nutrition plan was the same for all the patients and is described in table 1.

Table 1. Daily carbohydrate consumption for the in-silico patients

Time	7a.m.	10a.m.	12a.m.	3p.m.	6p.m.	10p.m.
Intake [g]	30	10	45	10	60	10

Table 2. Selected model for each patient

	Subject										
	Avg	**#1**	**#2**	**#3**	**#4**	**#5**	**#6**	**#7**	**#8**	**#9**	**#10**
n_a	4	4	4	6	6	4	4	4	4	4	4
n_b	8	6	10	10	10	6	10	10	10	10	10
n_c	10	9	12	6	10	6	8	12	10	10	10
n_d	6	9	10	12	8	9	8	6	12	12	12
n	150	250	250	350	250	250	350	250	250	250	250

Table 3. MPC settings

Subject	N	Q	R_{basal}	R_{bolus}	ρ_{hypo}	ρ_{hyper}	n_σ
#1, #3-10	12	100	0.5	0.1	10^5	10^4	2
#2, Avg	12	100	0.1	0.1	10^5	10^4	2

4.1 Data collection and model identification

In order to identify high-quality GPs for each patient, it was necessary to collect large amounts of data about all the conditions a diabetic subject can run into. For this reason, data about each patient were collected over a 35 days simulation with the goal to lead the patient to hyperglycaemia, euglycaemia and hypoglycaemia by randomly changing the basal insulin provided over the simulation time and giving random boluses at mealtime.

Giving random values of basal insulin and boluses requires knowing an upper bound for the two, namely the parameters $\overline{u}_{\text{basal}}$ and $\overline{\overline{u}}_{\text{bolus}}$ introduced in (8). In order to find these values it was necessary to run multiple simulations on each patient and to fix the upper bounds in order to avoid too dangerous hypoglycaemia and hyperglycaemia events.

Once the dataset $\mathcal{D}$ was built, the process explained in Section 3.1 was followed to determine n_a, n_b, n_c, n_d and n and find the best GP for each patient. In particular, in order to limit the number of parameter combinations to be evaluated, the following assumptions were made:

$$n_a \in \{4, 6\}, \tag{9a}$$
$$n_b, n_c, n_d \in [6, 12], \tag{9b}$$
$$n \in \{150, 200, 250, 300, 350, 400\}. \tag{9c}$$

The parameters of the chosen models are shown in table 2.

4.2 Controller setting

The personalization of the here proposed controller to meet the needs of the single T1DM patient is not only guaranteed by the GP model that is fitted to the single subject's data, but also by the possibility to choose different settings for the zone MPC. In particular, the parameters that can be changed to optimize the performance of the controller on the single subject are N, Q, R_{basal}, R_{bolus}, ρ_{hypo}, ρ_{hyper} and n_σ.

The settings that were chosen for the in-silico patients after several trials were similar. In particular, only R_{basal} took different values depending on the patient. The parameters for all the patients are presented in table 3.

4.3 Results

The results of the simulations are promising; in particular, no patients ran into severe hypoglycaemia events and, apart from Adult 7, all the other subjects experienced a limited number of hyperglycaemia events.

To evaluate the performance of our controller on each patient, we used the Control-Variability Grid Analysis (CVGA) introduced in (Magni et al., 2008). The CVGA

is shown in figure 1 and proves the general quality of our controller, given that:

- 2 patients never reached hypoglycaemia and hyperglycaemia situations (A zone);
- 2 patients experienced some benign hyperglycaemia events and approached the euglycaemic range lower bound (B zone);
- 6 patients experienced some benign hyperglycaemia events and never approached the euglycaemic range lower bound (Upper B zone);
- One patient, in particular Adult 7, ran into some severe hyperglycaemia events, but never approached hypoglycaemic levels (Upper C zone).

Figure 2 displays the evolution of the blood glucose level over the three-day simulation for three representative patients. In detail, it is possible to see how the controller was able to keep Adults 5 and 7 almost always within the euglycaemic range, responding promptly to carbohydrate intakes with well-computed boluses, while being too conservative in the case of Adult 7, who experienced several long-lasting hyperglycaemia events due to an insufficient infusion of basal insulin and small boluses. Other common

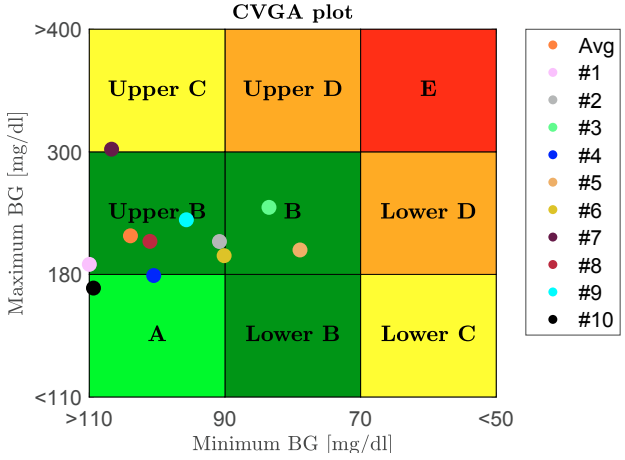

Fig. 1. CVGA of the simulations run for the 11 adult patients. Adult 7 was the only one reaching noticeable hyperglycaemia values, while the other patients were treated successfully with our MPC.

performance indices are shown in tables 4. Note how almost all the patients spent between the 85% and the 100% of the simulation time in euglycaemic range, the only exception being Adult 7. Furthermore, none of the patients ran into hypoglycaemia events.

In addition to achieving encouraging results, the proposed MPC always had execution times that were consistent with the chosen sampling time. Indeed, as shown in figure 3, only the execution time of the controller built for adult 6 approached the sampling time, but never exceeded it.

5. CONCLUSION

We proposed a zone MPC based on GPs for the control of the blood glucose level in T1DM patients and tested it on a total of 11 in-silico adult patients on the FDA-approved UVA/Padova simulator.
The main advantages of the introduced control method reside in the possibility to customize it to the individual

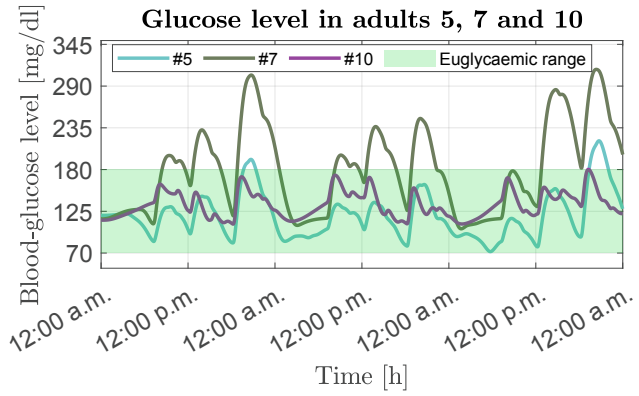

Fig. 2. Evolution of the blood glucose level in three representative patients: adult 7 ran into severe hyperglycaemia events, adult 5 experienced rare hyperglycaemia events, while adult 10 stayed within the euglycaemic range.

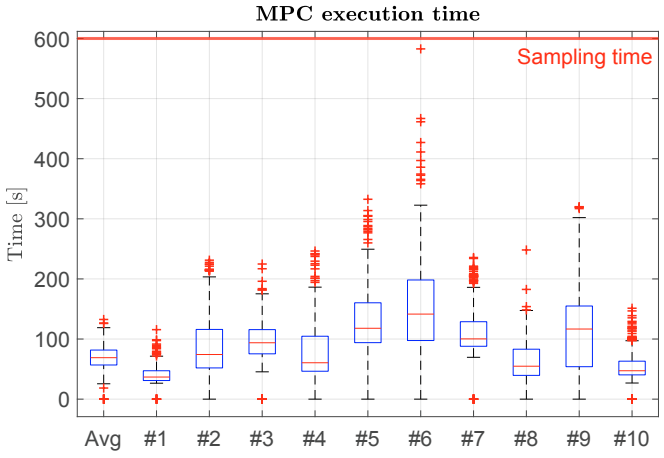

Fig. 3. Box plots of the single MPC step execution times related to each patient. The execution time of a single MPC step was never greater than the sampling time.

T1DM patient and its ability to compute the amounts of basal insulin and insulin boluses that have to be injected automatically and without exploiting any first-principle models of the endocrine-metabolic system. Moreover, modelling the system as a GP allows us to exploit the prediction uncertainties provided by GPR to avoid situations that are likely to bring the treated patient to hypoglycaemia and hyperglycaemia events.
The results obtained on the UVA/Padova simulator are encouraging and show how most of the in-silico patients benefited from the usage of our MPC. In particular, none of them experienced hypoglycaemia events. The CVGA shown in Figure 1 and the performance indices presented in table 4 clearly show that our controller tends to avoid hypoglycaemia more than it avoids hyperglycaemia, which is consistent with the MPC settings that we chose, specifically $\rho_{\text{hypo}} > \rho_{\text{hyper}}$.
The most critical issue about our control method is the collection of the data that are needed to fit a GP that models well the behavior of a subject's metabolism. Indeed, it is necessary to bring the patient to multiple hypoglycaemia and hyperglycaemia events in order for the identified GP to make good predictions in all situations. However, we are confident that the usage of constrained GPs, namely GPs that are trained in a way that allows one to incorporate

Table 4. Performance metrics of the closed-loop system on the 11 in-silico patients

Subject	**Avg**	**#1**	**#2**	**#3**	**#4**	**#5**	**#6**	**#7**	**#8**	**#9**	**#10**
Mean glucose [mg/dl]	141.55	140.17	134.43	130.38	131.24	119.66	122.93	181.05	145.10	140.49	133.68
Glucose SD [mg/dl]	28.13	17.78	24.67	33.61	20.48	30.17	24.07	57.81	26.79	35.74	17.26
Time in [70,140]mg/dl [%]	57.76	55.45	71.49	71.42	66.30	76.28	80.84	33.28	56.79	59.52	63.13
Time in [70,180]mg/dl [%]	89.93	96.41	94.38	93.57	98.10	93.91	95.76	51.15	88.43	85.81	99.79
Time above 180mg/dl [%]	10.07	3.59	5.62	6.43	1.90	6.09	4.24	48.85	11.57	14.19	0.21
Time below 70mg/dl [%]	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

a priori information about the underlying function, could lead to good results even by collecting fewer data and avoiding hypoglycaemia and hyperglycaemia events.

REFERENCES

Abuin, P., Rivadeneira, P.S., Ferramosca, A., and González, A.H. (2020). Artificial pancreas under stable pulsatile MPC: Improving the closed-loop performance. *Journal of Process Control*, 92, 246–260.

Bahremand, S., Ko, H.S., Balouchzadeh, R., Felix Lee, H., Park, S., and Kwon, G. (2019). Neural network-based model predictive control for type 1 diabetic rats on artificial pancreas system. *Medical & biological engineering & computing*, 57(1), 177–191.

Boiroux, D., Bátora, V., Hagdrup, M., Wendt, S.L., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2018). Adaptive model predictive control for a dual-hormone artificial pancreas. *Journal of Process Control*, 68, 105–117.

De Paula, M., Acosta, G.G., and Martínez, E.C. (2015). On-line policy learning and adaptation for real-time personalization of an artificial pancreas. *Expert Systems with Applications*, 42(4), 2234–2255.

Del Favero, S., Toffanin, C., Magni, L., and Cobelli, C. (2019). Deployment of modular MPC for type 1 diabetes control: the Italian experience 2008–2016. In *The Artificial Pancreas*, 153–182. Elsevier.

Dutta, S., Kushner, T., and Sankaranarayanan, S. (2018). Robust data-driven control of artificial pancreas systems using neural networks. In *International Conference on Computational Methods in Systems Biology*, 183–202. Springer.

Gondhalekar, R., Dassau, E., and Doyle III, F.J. (2016). Periodic zone-MPC with asymmetric costs for outpatient-ready safety of an artificial pancreas to treat type 1 diabetes. *Automatica*, 71, 237–246.

González, A.H., Rivadeneira, P.S., Ferramosca, A., Magdelaine, N., and Moog, C.H. (2017). Impulsive zone mpc for type i diabetic patients based on a long-term model. *IFAC-PapersOnLine*, 50(1), 14729–14734.

González, A.H., Rivadeneira, P.S., Ferramosca, A., Magdelaine, N., and Moog, C.H. (2020). Stable impulsive zone model predictive control for type 1 diabetic patients based on a long-term model. *Optimal Control Applications and Methods*, 41(6), 2115–2136.

Hajizadeh, I., Rashid, M., and Cinar, A. (2019). Plasma-insulin-cognizant adaptive model predictive control for artificial pancreas systems. *Journal of process control*, 77, 97–113.

Hovorka, R., Canonico, V., Chassin, L.J., Haueter, U., Massi-Benedetti, M., Federici, M.O., Pieber, T.R., Schaller, H.C., Schaupp, L., Vering, T., et al. (2004). Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes. *Physiological measurement*, 25(4), 905.

Huyett, L.M., Dassau, E., Zisser, H.C., and Doyle III, F.J. (2015). Design and evaluation of a robust PID controller for a fully implantable artificial pancreas. *Industrial & engineering chemistry research*, 54(42), 10311–10321.

Kovatchev, B. (2018). Automated closed-loop control of diabetes: the artificial pancreas. *Bioelectronic Medicine*, 4(1), 1–12.

Magni, L., Raimondo, D.M., Man, C.D., Breton, M., Patek, S., De Nicolao, G., Cobelli, C., and Kovatchev, B.P. (2008). Evaluating the efficacy of closed-loop glucose regulation via control-variability grid analysis. *Journal of diabetes science and technology*, 2(4), 630–635.

Moon, S.J., Jung, I., and Park, C.Y. (2021). Current advances of artificial pancreas systems: a comprehensive review of the clinical evidence. *Diabetes & Metabolism Journal*, 45(6), 813–839.

Ortmann, L., Shi, D., Dassau, E., Doyle, F.J., Leonhardt, S., and Misgeld, B.J. (2017). Gaussian process-based model predictive control of blood glucose for patients with type 1 diabetes mellitus. In *2017 11th Asian Control Conference (ASCC)*, 1092–1097. IEEE.

Patra, A.K. and Rout, P.K. (2017). Adaptive sliding mode Gaussian controller for artificial pancreas in TIDM patient. *Journal of Process Control*, 59, 13–27.

Pillonetto, G. (2018). System identification using kernel-based regularization: New insights on stability and consistency issues. *Automatica*, 93, 321–332.

Quinonero-Candela, J., Girard, A., and Rasmussen, C.E. (2003). Prediction at an uncertain input for Gaussian processes and relevance vector machines-application to multiple-step ahead time-series forecasting.

Rasmussen, C.E. (2003). Gaussian processes in machine learning. In *Summer school on machine learning*, 63–71. Springer.

Rosales, N., De Battista, H., and Garelli, F. (2022). Hypoglycemia prevention: PID-type controller adaptation for glucose rate limiting in Artificial Pancreas System. *Biomedical Signal Processing and Control*, 71, 103106.

Shi, D., Dassau, E., and Doyle, F.J. (2018). Adaptive zone model predictive control of artificial pancreas based on glucose-and velocity-dependent control penalties. *IEEE Transactions on Biomedical Engineering*, 66(4), 1045–1054.

Soru, P., De Nicolao, G., Toffanin, C., Dalla Man, C., Cobelli, C., Magni, L., Consortium, A.H., et al. (2012). MPC based artificial pancreas: strategies for individualization and meal compensation. *Annual Reviews in Control*, 36(1), 118–128.

The Epsilon Group, U.o.V. (2016). DMMS.R (Version 1.1) [Software]. Retrieved from https://tegvirginia.com/.

Toffanin, C., Messori, M., Di Palma, F., De Nicolao, G., Cobelli, C., and Magni, L. (2013). Artificial pancreas: model predictive control design from clinical experience.