OPTIMAL CONTROL APPLICATIONS AND METHODS
Optim. Control Appl. Meth. (2015)
Published online in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/oca.2220

A grid-based tool for optimal performance monitoring of a glycemic regulator

Luis Ávila and Ernesto Martínez*,†

INGAR (CONICET-UTN), Avellaneda 3657, Santa Fe, GJC S3002, Argentina

SUMMARY

Recent technology breakthroughs towards a fully automated artificial pancreas give rise to the need of new monitoring tools aiming at increasing both reliability and performance of a closed-loop glycemic regulator. Based on error grid analysis, an insightful monitoring tool is proposed to assess if a given closed-loop implementation respects its specification of an optimally performing glycemic regulator under uncertainty. The optimal behavior specification is obtained using linearly solvable Markov decision processes, whereby the Bellman optimality equation is made linear through an exponential transformation that allows obtaining the optimal control policy in an explicit form. The specification for the desired glucose dynamics is learned using Gaussian processes for state transitions in an optimally performing artificial pancreas. By means of the proposed grid, the specification is *vis-à-vis* compared with glucose sensor readings so that any significant deviation from the expected closed-loop performance under abnormal or faulty scenarios can be detected.

Received 4 July 2014; Revised 15 September 2015; Accepted 19 October 2015

KEY WORDS: artificial pancreas; optimal control applications; performance monitoring; stochastic optimal control; type 1 diabetes; uncertain systems

1. INTRODUCTION

Type 1 diabetes is caused by selective destruction of the pancreas β-cells, which results in deficient insulin production for glucose regulation. Thus, to restore normal glucose levels, exogenous insulin delivery is required. Because it is difficult to achieve normoglycemia with intermittent insulin dosage, continuous intervention of healthcare professionals is required for proper diabetes management. As a result, the majority of people with diabetes cannot meet their treatment goals [1]. Performance assessment of glycemic regulation is a crucial issue in diabetes care. Reliable information about blood glucose variability allows clinicians and patients to evaluate the efficacy of insulin infusion regimes to maintain glucose levels within the normoglycemic range. Monitoring the variability of glucose levels also allows adjustment in diet, exercise, and medications to improve blood glucose regulation. Furthermore, monitoring glycemic variability helps in diminishing the risks of hypoglycemia or hyperglycemia events and long-term complications [2]. However, the dynamics of human physiology, errors in glucose sensors, and failures in insulin infusion pumps give rise to a number of challenges for implementing an artificial pancreas (AP). Increasing scientific and industrial effort is focused on the development of automatic regulation systems to control insulin delivery in people with diabetes [3]. The ultimate goal in closed-loop control of glycemia is not just finding the optimal insulin rates that can effectively reduce high blood glucose levels but also infusing it in such a way that the blood glucose level can mimic the body's natural regulatory mechanism [4].

*Correspondence to: Ernesto Martinez, INGAR (CONICET-UTN), Avellaneda 3657, Santa Fe, GJC S3002, Argentina.
†E-mail: ecmarti@santafe-conicet.gov.ar

The safety-critical condition of such an automated device makes its performance monitoring task of paramount importance in order to reduce glucose level variability and to minimize the risks of dangerous excursions outside the normoglycemic range. Hopefully, closing the glycemic control loop with a fully automated system will certainly improve the quality of life for insulin-dependent patients. Such a device is made up of a glucose measuring device, an automated insulin infusion pump, and a feedback control strategy that calculates the insulin infusion rate based on an error-prone glucose signal. Obviously that, to guarantee an optimal glucose regulation, safe implementation requires proper functioning of each individual component at all times. Because human body is a highly nonlinear, robust, and adaptive physiological system, there exists a close relationship between stochastic optimal control and diabetes care.

The widespread availability of user-friendly monitoring strategies for glucose control is a key issue to the acceptance of an AP. The criteria to evaluate the efficacy of a monitoring tool should be based on the usefulness of the test results to improve glycemic control [5], so that the acceptance of such system by patients and clinicians can be significantly increased. Continuous glucose monitoring will certainly improve daily diabetes treatment, but the realization of this promise awaits a complementary shift in the way that glycemic data are processed [6]. Similarly to other regulatory loops in physiological systems, accessibility to some key variables requires a highly invasive mechanism, which is not always feasible [7]. Moreover, periodical measuring of plasma glucose to check if the AP is performing correctly results impracticable for ambulatory patients, particularly at night time [8]. To solve this problem, a specification of the correct behavior of an AP under uncertainty is proposed, and then it is compared with samples obtained through a subcutaneous sensor. This allows contrasting the obtained sensor readings through a *vis-à-vis* comparison with the desired glycemic behavior, without repetitively solving a computationally demanding algorithm and without the on-going active intervention of a clinician. This also has the advantage of making the monitoring task simple enough to be executed on a mobile device.

Even though the specified behavior describes the optimal manner the AP is expected to work, this only holds true if a number of factors causing uncertainty are taken into account in specifying the glucose regulator behavior. For instance, the precision of the insulin pump, physiological and physical lag times, and the calibration error of the glucose sensor must be accounted for. Even though there exists a vast literature on using optimal control methods for designing a glucose regulator, only a few research works focus on the problem of performance monitoring of an AP. Undoubtedly, the most significant development in control performance monitoring was due to the work of Harris [9], who proposed the concept of a minimum variance controller as the characteristic behavior of an optimal regulator. Thus, performance monitoring can be based on comparing the observed mean squared error of the controlled output with its corresponding minimum variance. This strategy has already been implemented in monitoring the performance of a control loop for a diabetic patient [10]. However, because of unavoidable sources of uncertainty and glycemic variability, performance assessment tools using the minimum variance method are unable to differentiate normal from abnormal behavior of an AP [11, 12]. Moreover, for performance monitoring, it is usually assumed that the control system cannot influence or modify the environment in which it operates. But an AP must perform in a time-varying and highly nonlinear environment [13], namely, the human body, which defines a feedback loop with a glycemic regulator. As the AP is a situated entity interacting with a highly regulated system, its behavior cannot be monitored using a fixed reference, threshold, or bound.

This work highlights that stochastic optimal control is the cornerstone for defining the specification of an AP under uncertainty. To this aim, a probabilistic characterization of the expected glycemic behavior is built upon a stochastic process specification of the glucose regulator. The desired optimal behavior is obtained analytically using a class of Markov decision processes that are linearly solvable [14, 15]. Through an exponential transformation, the Bellman equation for such problems can be made linear despite nonlinearity in the stochastic dynamical models, which facilitates applying efficient numerical methods. The AP specification is modeled by a prior Gaussian distribution for state transitions, which accounts for different sources of glycemic variability. The availability of an optimal control policy allows simulating the desired behavior over time and comparing it with its specification in order to identify deviations from the desired performance.

DOI: 10.1002/oca

To this aim, an offline clinical monitoring tool based on error grid analysis (EGA) is proposed to detect failures and performance degradation in the closed-loop glycemic regulator. The major advantage for the closed-loop error grid analysis (CL-EGA) is that it can be employed for visually detecting significant deviations of an AP from its specified behavior corresponding to an optimally performing glycemic regulator.

2. OPTIMAL CONTROL UNDER UNCERTAINTY

2.1. Optimal control policy

To achieve a proper characterization of the optimal behavior of a glycemic regulator under uncertainty, the novel approach of linearly solvable Markov decision process (LSMDP) is used [14], whereby the Bellman optimality equation is made linear through an exponential transformation that allows obtaining the optimal control policy in an explicit form.

Consider a controller choosing actions over time (the glycemic regulator), an uncertain dynamical system whose state is affected by those actions (the body), and a performance criterion that the controller seeks to achieve (normoglycemia). We employ a structured formulation that greatly simplifies the construction of optimal control laws in continuous domains, because an exhaustive search over actions is avoided and the problem becomes linear. To this aim, the optimal cost-to-go function $v(\mathbf{x})$ is defined as the expected cumulative cost for starting at state $\mathbf{x}$ and acting optimally thereafter. By acting on the uncertain glycemic dynamics, the controller generates a sequence of actions that optimize a performance criterion in the long run. As a result, if the regulator performs away from the specified behavior, it must pay a price for suboptimal reshaping of the glycemic dynamics [16]. Any increase in the price to be paid is an evidence of a performance loss. Hence, $v(\mathbf{x})$ equals the minimum of an immediate cost $\ell(\mathbf{x}, u)$ plus an expected cost-to-go $\mathrm{x}'\sim\mathrm{u}(\mathrm{x}'|\mathrm{x})$ at the next state $\mathbf{x}'$ (see (4) in the succeeding texts). The subscript indicates that the expectation is taken with respect to the state transition probability distribution $p^u(\mathbf{x}'|\mathbf{x}, u)$, which is induced by the control action u.

A special feature of this formulation is that, instead of asking the controller to specify symbolic actions that are later replaced with transition probabilities, we allow the controller to specify transition probabilities directly. Let $p^u(\mathbf{x}'|\mathbf{x}, u)\equiv u(\mathbf{x}'|\mathbf{x})$ denote the transition probability given a control policy $u=\pi(\mathbf{x})$, whereas $p^0(\mathbf{x}'|\mathbf{x})$ models the transition probability for the uncontrolled or passive glycemic dynamics. Regarding uncertainty, the state transition function in this problem obeys to a controlled diffusion process with the following form:

$$d\mathbf{x} = \mathrm{a}(\mathbf{x})dt + B(\mathbf{x})(udt + \sigma d\omega) \tag{1}$$

where $\omega\in\mathbb{R}^{1_u}$ and σ denote a Brownian noise and its scaling parameter, respectively. The expression a $(\mathbf{x})$ describes the passive dynamics, and $B(\mathbf{x})$ is the input–gain matrix.

In order to express (1) in a more convenient form, the h-step transition probability for the passive or uncontrolled dynamics p^0 is expressed as a Gaussian distribution $\mathcal{N}$ as

$$p^0\left(\mathbf{x}'\middle|\mathbf{x}\right) = N\left(\mathbf{x}'\middle|\mathbf{x} + h\mathrm{a}(\mathbf{x}) + hB(\mathbf{x}), h\sigma B(\mathbf{x})^T B(\mathbf{x})\right) \tag{2}$$

The controlled diffusion process p^u is approximated as a deterministic function expressed as a Gaussian distribution whose mean and covariance are given as

$$p^u\left(\mathbf{x}'|\mathbf{x}\right) = N\left(\mathbf{x}'|\mathbf{x} + h(\mathrm{a}(\mathbf{x}) + hB(\mathbf{x})u), \sum\right) \tag{3}$$

One way of thinking the net effect of control actions is noting how they change the distribution of the next state from $(\mathbf{x}+h\mathrm{a}(\mathbf{x})+hB(\mathbf{x}),\sum)$ to $(\mathbf{x}+h(\mathrm{a}(\mathbf{x})+hB(\mathbf{x})u),\sum)$, where $\sum=h\sigma B(\mathbf{x})^T B(\mathbf{x})$ is the

DOI: 10.1002/oca

covariance. In other words, the controller shifts the probability distribution from one region of the state space to another [16]. Specifically, it is assumed that a diabetic patient has an uncontrolled glucose dynamics that gives rise to a distribution p^0 over future states; then, the regulator acts by modifying this distribution obtaining a controlled glucose dynamics p^u.

2.2. Optimal control policy

An insulin control policy $\pi(\mathbf{x})$ is thus defined as a probability of selecting the action u, which corresponds to a change to the insulin infusion rate, at the glycemic state $\mathbf{x}$. The main objective is to find an optimal insulin infusion policy $\pi^*(\mathbf{x})$ that minimizes the expected cumulative cost function $v(\mathbf{x})$ as

$$v^*(\mathbf{x}) = \min_u \left\{ \ell(\mathbf{x}, \pi(\mathbf{x})) + \mathbf{E}_{\mathbf{x}' \sim u(\mathbf{x}'|\mathbf{x})} \left[v\left(\mathbf{x}'\right) \right] \right\} \tag{4}$$

where $\mathbf{x}'$ denotes the next glycemic state for a given action u. The minimum cost-to-go for starting at a state $\mathbf{x}$ and acting optimally thereafter enables greedy computation of optimal actions. Equation (4) is fundamental to optimal control theory and is called the Bellman optimality equation. The Bellman equation can be simplified by assuming that the immediate cost function has the following form:

$$\ell(\mathbf{x}, u) = hq(\mathbf{x}) + KL\left(u\left(\mathbf{x}'|\mathbf{x}\right) \,||\, p\left(\mathbf{x}'|\mathbf{x}\right)\right) \tag{5}$$

Here, the state cost $q(\mathbf{x})$ is an arbitrary function encoding how desirable different glucose levels are. Because the controller acts directly over the transition probabilities, it makes sense to measure its magnitude in terms of the difference between the controlled transition probability and the uncontrolled transition probability by means of the Kullback–Leibler divergence. The distance can be understood as the price to be paid for the optimal shifting by the insulin infusion action u of the passive dynamics corresponding to an uncontrolled patient. An exponential function of the value function $v(\mathbf{x})$ is used to highlight that the most desired states are those where the cost-to-go is small. Even though the Bellman equation is nonlinear in terms of the value function, it yields a linear equation in a new variable $z(\mathbf{x})$ when using the exponential transformation

$$z(\mathbf{x}) = \exp(-v^*(\mathbf{x})) \tag{6}$$

and combining (5) and (6), the Bellman equation can be expressed as a linear function

$$-\log(z(\mathbf{x})) = q(\mathbf{x}) \min_u \left\{ \mathbf{E}_{\mathbf{x}' \sim u(\mathbf{x}'\mathbf{x})} \left[\frac{u\left(\mathbf{x}'|\mathbf{x}\right)}{p(\mathbf{x}'|\mathbf{x}) z(\mathbf{x}')} \right] \right\} \tag{7}$$

The last expression resembles the Kullback–Leibler divergence between $u(\mathbf{x}'|\mathbf{x})$ and $p(\mathbf{x}'|\mathbf{x})\ z(\mathbf{x}')$, but it should be first normalized by introducing an integral operator $G[z](\mathbf{x})$ defined as

$$G[z](\mathbf{x}) = \int p^0\left(\mathbf{x}'|\mathbf{x}\right) z\left(\mathbf{x}'\right) d\mathbf{x}' \tag{8}$$

The use of this integral operator allows (7) to be rewritten as

$$z(\mathbf{x}) = \exp(-hq(\mathbf{x})) G[z](\mathbf{x}) \tag{9}$$

which is known as the *desirability* function [15].

In contrast to the cost function $v(\mathbf{x})$, the negative exponential portrays that states are more desirable. Once the desirability function is found, the optimal control policy is computed using (6) as

DOI: 10.1002/oca

$$u^*(\mathbf{x}) = -\sigma^2 B(\mathbf{x})^{\mathrm{T}} v_{\mathbf{x}}(\mathbf{x}) \tag{10}$$

In this manner, optimal action selection can be expressed analytically given the optimal cost-to-go. Thus, instead of finding a trajectory-based solution, the goal is to find a global optimal policy over the entire state space.

The continuous problem given in (1) can now be solved by choosing a set of states $\{\mathbf{x}_n\}$ and adjusting the matrix P of transition probabilities from state $\mathbf{x}$ to the next $\mathbf{x}'$ given by the passive dynamics described in (2). The min operator in (4) can be dropped, and the Bellman equation is then expressed in terms of the transformed variable $z(\mathbf{x})$. As the integral operator $G[z]$ $(\mathbf{x})$ is linear, thus, (9) is also linear in $z(\mathbf{x})$ and can be expressed in vector notation. Defining the vector $\boldsymbol{z}$ with elements $z(\mathbf{x}_n)$ and the matrix Q with elements $\exp(-hq(\mathbf{x}_n))$ on its main diagonal, we obtain

$$\boldsymbol{z} = QP\boldsymbol{z} \tag{11}$$

This expression has the same functional form that an eigenvector problem has, and then can be solved iteratively by taking advantage of the exponential form of the value function [14].

2.3. Modeling glycemic dynamics

Diabetes mellitus is characterized by the inability of the pancreas to properly control blood glucose concentration. Current subcutaneous treatments often result in poor maintenance of normoglycemia (glucose levels within 80–140 $\mathrm{mg\,dl^{-1}}$), and excessive variation in blood glucose levels is observed because of the inefficacy of intermittent control. In order to achieve a proper characterization of an optimal regulator behavior, glycemic variability is simulated by means of a stochastic processes superimposed on an otherwise deterministic model of the glucose–insulin dynamics. The Lehmann and Deutsch model [17] is used as the basis to describe such deterministic dynamics. By adding an Ito's process to the model, glycemic variability in a type 1 diabetic patient is obtained as

$$dBG = \left(\frac{G_{in} + NHGB - G_{out} - G_{ren}}{V_G}\right)dt + \sigma d\omega \tag{12}$$

where BG is the plasma glucose concentration, G_{in} is the systemic appearance of glucose via glucose absorption from the gut, $NHGB$ is the net hepatic glucose balance, G_{out} is the overall rate of peripheral and insulin-dependent glucose utilization, G_{ren} is the renal excretion of glucose, and V_G is the glucose distribution volume. A value of $\sigma = 0.25$ was suggested to be representative of glycemic variability in the work of Acikgoz and Diwekar [18]. Such value allows considering daily variations in glucose concentration on a diabetic patient, also including different unknown sources of uncertainty such as errors in the model, inaccurate measuring, and unpredictable behavior of the metabolism in diabetic patients.

Modeling glycemic variability through a diffusion process, a cohort of *in silico* subjects that accounts sufficiently well for the observed inter-subject and intra-subject variability is obtained. For all simulations hereafter, the meal intake regime in Table I is used.

Error-prone measurements of BG concentration are obtained by a subcutaneous glucose sensor, which is only an estimation of the actual plasma glucose concentration. As the sensor needle is placed in the subcutaneous tissue, it determines the interstitial fluid (IG) concentration instead of

Table I. Carbohydrate intake schedule.

Carbohydrate content [g]	47	16	63	31	63	31
Meal times [h]	8.00	10.00	12.30	16.00	19.30	22.00

DOI: 10.1002/oca

the plasma glucose directly. To this aim, the sensor response is described in (13). Each *IG* value is obtained by integrating a *BG-IG* dynamics, where ρ is the static gain (considered equal to 1) and τ is the time lag constant

$$\begin{aligned} dIG(k) &= -\frac{1}{\tau}IG(k) + \frac{\rho}{\tau}BG(k) \\ s(k) &= (1+\xi(k))IG(k) + \gamma(k) \end{aligned} \tag{13}$$

As a result, sensor readings $s(k)$ are corrupted by a random time-varying calibration error plus a white Gaussian noise process $\gamma(k)$. It is worth noting that the expression *BG* instead of *IG* is used hereafter, but recall that it refers to the outcome of a continuous glucose monitor. In Figure 1, plasma glucose levels and sensor responses are depicted for different sensor calibration errors and time lag values using the optimal control policy (see Section 3.3, for details).

2.4. *Glucose-insulin passive dynamics*

As mentioned, the passive dynamics represents the behavior of the stochastic dynamics in the absence of control actions. Likewise, in a type 1 diabetic patient, the passive dynamics represents the behavior of *BG* levels when no exogenous insulin is administrated. Because the ability to mitigate glycemic variability has been lost, the patient may experiment hyperglycemic or hypoglycemic events as *BG* values vary without control over a wide range. As a consequence, this passive dynamics makes possible for the learning controller to visit a large number of glycemic states, which are essential for obtaining an optimal control policy. The passive dynamics is then crucial in order to specify the desired behavior of an AP, yet it is unknown in general and needs to be somehow approximated using a model. To begin with, notice that by discretizing the dynamical system in (1), the *h*-step state transition probability for the uncontrolled glucose dynamics can be expressed as a Gaussian distribution as given in (2). Then, the continuous problem is approximated in discrete time by choosing a set of *BG* states $\{\mathbf{x}_n\}$ and adjusting the matrix $P_{k,k+1}$ of transition probabilities from $\mathbf{x}_k$ to $\mathbf{x}_{k+1}$ given by the passive dynamics distribution.

Because the passive dynamics accounts for the space of all possible state transitions in an uncontrolled patient, it is possible to estimate the passive dynamics by using a reduced version of the model described in (12). In fact, the Lehmann and Deutsch model previously described is an augmented representation of the two-compartment minimal model described in Bergman *et al.* [19]

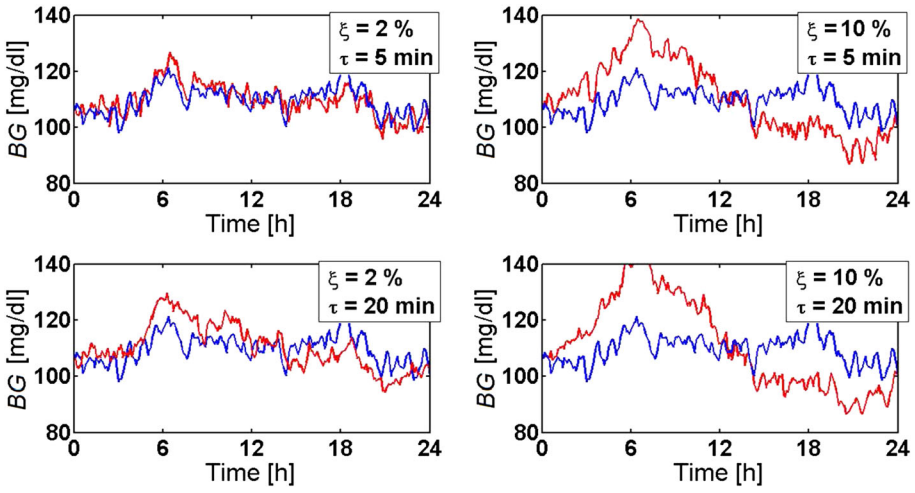

Figure 1. Sensor reading errors distorting plasma glucose levels because of miscalibration and excessive time lags.

DOI: 10.1002/oca

$$\begin{aligned} \frac{dBG}{dt} &= (p_1 - Ia)BG + p_1 G_b \\ \frac{dIa}{dt} &= p_2 Ia + p_3 I(t) \end{aligned} \tag{14}$$

where Ia represents the time course or profile for the insulin infusions and G_b is the basal glucose level. In this minimal model of the glucose kinetics, the insulin action $I(t)$ enters at a remote compartment with the goal of speeding up the glucose disappearance process and at the same time inhibits every hepatic glucose production. The model parameters p_i are given in Table II. Despite the limitations of the Bergman's minimal model, it can properly describe glucose–insulin state transitions in a diabetic patient.

The glycemic regulator is thus based on two physiological variables, which can be readily known. The control task is mainly concerned with reducing glycemic variability while minimizing the deviation from the desirable BG levels. For this purpose, the cost function is conveniently designed in such a way it guarantees an acceptable behavior of blood glucose dynamics within a target band (80–140 $\mathrm{mg\,dl^{-1}}$). The state cost function $q(\mathbf{x})$ is thus represented by a square exponential function that enforces normoglycemia, whereas it penalizes large deviations from the desired BG levels, that is, values as close as possible to $BG = 110\,\mathrm{mg\,dl^{-1}}$ and $Ia = 30\,\mathrm{mU\,l^{-1}}$. Note that this is a two-dimensional state problem with $\mathbf{x}_k = [BG_k, Ia_k]$ and control action $u_k = \Delta \mathrm{I}_k$, which accounts for the amount of plasma insulin entering into the remote compartment. A restriction is imposed so that the rate of change of the control action is not higher than a preset value as follows given as $\Delta \mathrm{I}_k \leq 8\,\mathrm{mU\,min^{-1}}$.

To approximate the continuous problem in (12), we use a state space discretization by means of a 151-by-151 grid over the intervals $BG \in [0, 220]\ \mathrm{mg\,dl^{-1}}$ and $Ia \in [0, 60]\ \mathrm{mU\,l^{-1}}$. The passive dynamics is then constructed by discretizing the time axis, using a time step $h = 0.05$, and by defining probabilistic transitions among discrete states as in (2) such that the mean and variance of the continuous state dynamics are preserved. The noise distribution is discretized at 9 points spanning ±3 standard deviations and using a noise scale parameter $\sigma = 0.1$. A small value of σ is chosen here to represent glycemic variability because inaccurate measurements are included in the sensor error.

Once the passive dynamics matrix P and the cost matrix Q are found, the desirability function can be optimized by means of (11). One way to find the solution is the power iteration method – which in the present context is equivalent to value iteration in the exponentiated form. Despite this equivalence, however, solving (11) can be orders of magnitude faster than value iteration for generic MDPs, because the optimal controls are found analytically and the value iteration can converge slowly, whereas using the exponential transformation, convergence is linear. The optimal control policy was subsequently derived from the obtained desirability function using (6) and (11).

Results are displayed in Figure 2, where the axes denote the variables used to describe the glycemic state of a patient. Figure 2(a) depicts the scheme of the Bergman's minimal model used to approximate the passive dynamics. Even if this reduced model is used to approximate the passive dynamics, it is worth noting that the optimal control policy will be applied in the augmented model presented in (12) including the sensor signal response. Figure 2(b) graphically depicts the state cost function $q(\mathbf{x})$ centered on the desired basal values for BG and Ia. In Figure 2(a), blue regions correspond to low values and red to high values. Figure 2(c) shows the optimal cost-to-go function obtained. As can be seen, states with higher costs are placed in the upper and lower boundaries of the plot.

Table II. Parameters of the Bergman's minimal model.

Parameter	Value
p_1	2.96×10^{-2} [$\mathrm{min^{-1}}$]
p_2	1.86×10^{-2} [$\mathrm{min^{-1}}$]
p_3	6.51×10^{-6} [$\mathrm{min^{-2}\,\mu U^{-1}\,ml^{-1}}$]
G_b	97 [$\mathrm{mg\,dl^{-1}}$]

DOI: 10.1002/oca

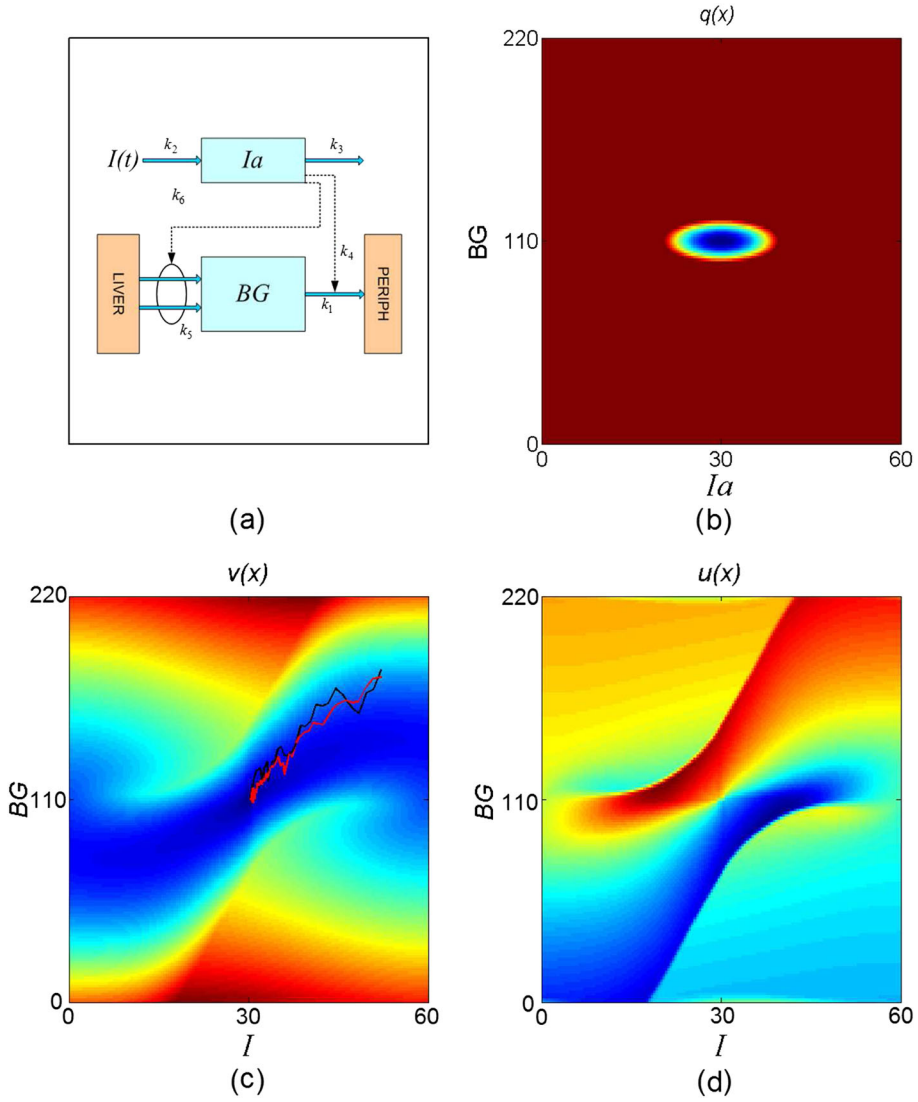

Figure 2. (a) Reduced model of the glucose-insulin dynamics. (b) State cost function $q(\mathbf{x})$. (c) The optimal cost-to-go function. Sample trajectories are generated using different values of the noise scale parameter σ. (d) The optimal control policy obtained.

The two small paths drawn over the optimal cost-to-go function correspond to stochastic trajectories generated by the optimal policy; the red one was obtained using a noise scale $\sigma = 0.1$, whereas the black trajectory corresponds to $\sigma = 0.25$. When *BG* is high, insulin infusion rates are also high so that costs can be lowered, whereas smaller insulin infusion rates are required for optimal control when *BG* is low. Finally, Figure 2(d) describes the obtained optimal control policy used as the specification for behavior monitoring of a glycemic regulator under uncertainty.

2.5. *Optimal behavior specification*

The key issue to be addressed in performance monitoring of a glycemic regulator is how a well-performing AP can be characterized in the face of uncertainty. The desired behavior of a given glucose–insulin system is mandatory to detect the discrepancies between the specification and degraded performance in a suboptimal implementation. To this aim, state transition probabilities are represented using fully probabilistic Gaussian processes (GP) models that provide information about confidence intervals for value function predictions and optimal actions.

DOI: 10.1002/oca

Gaussian processes [20] are stochastic processes used to describe distributions directly into the space of functions. It is implicit that training examples are informative for reliable predictions. When using GPs, it is assumed that the joint distribution of the data is a multivariate Gaussian. Consequently, the problem is to find a covariance function that explains the data properly. Interesting properties of GPs are providing a probabilistic approach to inductive modeling and giving uncertainty estimates through prediction variances. GP regression uses a collection of random variables to represent the value of the unknown function $f(\mathbf{x})$ for different inputs $\mathbf{x}$. A GP is fully specified by a zero mean function $m(\mathbf{x}_i)$ and a covariance function $\mathrm{k}(\mathbf{x}_i, \mathbf{x}_j)$, encoding correlations between pairs of random variables

$$\mathrm{k}(\mathbf{x}_i, \mathbf{x}_j) = \exp\left(-\gamma_r \left\|\mathbf{x}_i - \mathbf{x}_j\right\|^2\right) + \lambda_r \delta_{ij} \tag{15}$$

with $\gamma_r \geq 0$ the kernel width parameter, $\lambda_r \geq 0$ the noise variance, and $\mathbf{u}_k$ the Kronecker delta function. This prior for the kernel function constrains input samples that are nearby to have highly correlated outputs. Short-term transition dynamics are modeled based on interactions (real or simulated) with the glycemic regulatory system. Given any state vector $\mathbf{x}$, a separate GP model is trained for each state dimension x, in such a way the effect of uncertainty about its change due to a control action is modeled statistically as

$$\Delta x_k \sim \mathcal{GP}(m, \mathrm{k}) \tag{16}$$

where the training inputs to the model are the states, whereas the targets are the differences between the successor state and the state in which the action is applied.

We can build the optimal transition probability $u^*(\mathbf{x}_{k+1}|\mathbf{x}_k)$ as a GP model, namely, $\mathcal{GP}^*$, which describes the stochastic specification of the AP. On the other hand, a model $\mathcal{GP}^g$ describes any implemented glycemic regulator, modeled as the transition probability $u^g(\mathbf{x}_{k+1}|\mathbf{x}_k)$, which may deviate from optimal control under uncertainty. In order to detect abnormal or faulty conditions, the GP model used to characterize the glycemic regulator must be updated online. In Figure 3, a realization of the glucose stochastic process obtained by applying the optimal policy for a scale of variability $\sigma = 0.10$ is shown. To simulate an increase in glycemic variability in a diabetic patient, the noise scale parameter is changed to $\sigma = 0.50$ from the 12th hour onwards, while the optimal control policy is applied. In the lower part of Figure 3, the predicted state transition distributions are given: Shaded areas describe the uncertainty in predictions, whereas solid lines correspond to the prediction means. Note that an increase in glucose variability is shown not only for the predicted means but

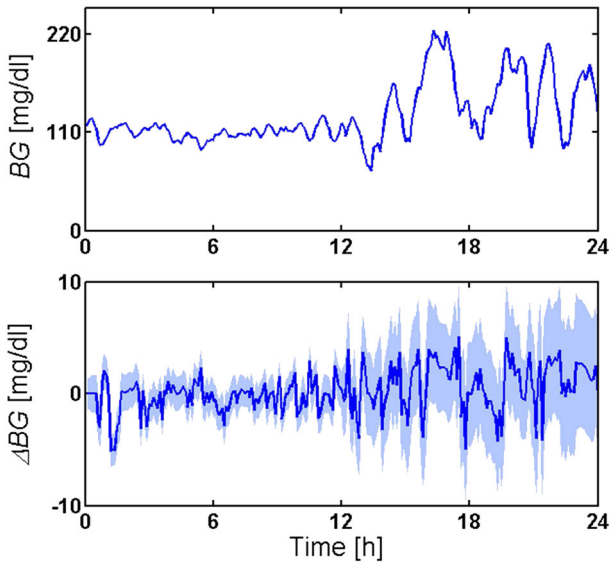

Figure 3. Effect of glycemic variability on the performance of a glycemic regulator. Mean and standard deviation values are generated through the noise scale parameter σ.

DOI: 10.1002/oca

also for the errors in the glucose level predictions, which reveals a significant degradation in the performance of the glycemic regulator because of a sudden increase in glycemic variability.

3. OPTIMAL CONTROL PERFORMANCE MONITORING

3.1. Implementation challenges of an artificial pancreas

Despite recent technology breakthroughs towards a fully automated AP, many issues still prevent safe and optimal operation of closed-loop therapies in diabetic patients. For example, continuous subcutaneous insulin infusion devices are prone to technical failures, being the obstruction of the catheter the most frequent event [21]. Moreover, sensor readings are affected not only by time lags and calibration errors – which increase with time – but also by environmental factors and other substances that can generate unstable output signals [22]. As the control algorithm may be executed in a mobile device, for example, a smart phone, wireless connection between the components of the glycemic regulator may be affected by delays or interference in data transmission. Even worse, ill-tuning of the control algorithm used to calculate the insulin bolus required can lead to severe hypoglycemia events. Thus, it is of paramount importance to address the issue of performance monitoring of an AP whose desired behavior should be established with respect to a formal specification of an optimally controlled glycemic dynamics. To this aim, the performance of a blood glucose regulator is evaluated in a range of treatment scenarios to detect any deviation of an implementation regarding its specification. It is important to note that we are dealing with a monitoring tool and not a diagnostic one. Thus, a warning about a deviant behavior does not necessarily need to differentiate between different probable causes, for example, a sensor calibration error or abnormal glycemic variability. Thus, monitoring is mostly concerned with converting glucose sensor data into warning signals accounting for any suspicious behavior of the AP and extracting meaningful information to prompt immediate action such as opening the glucose control loop [23].

3.2. Closed-loop error grid analysis

To further our concept of performance monitoring based on a specification of an optimally performing AP, let us introduce EGA methods designed to assess the accuracy of glucose sensors. EGA methods use a Cartesian diagram in which *BG* data measured through a glucose sensor are displayed in the y-axis against the *BG* data provided by the specified or reference method in the x-axis. For example, in Clarke's grid [24], the diagonal line represents the perfect agreement between the two measures, whereas points below and above the 45° line indicate overestimation and underestimation of the actual glycemic values, respectively. Another suitable tool to assess the efficacy of glucose variability control is the so-called control variability grid analysis proposed by Magni *et al.* [25]. The tool is appropriate for visualization of extreme glycemic excursions as well as detection of abnormal variability patterns. More closely related to our work, Chassin *et al.* [26] introduced a grading tool specifically designed to facilitate clinical assessment of closed-loop systems including that of glucose controllers. In similar works (refer to [27, 28]), the performance of glucose prediction algorithms has been tested using a reference measure.

Nevertheless, lacking an optimal control policy under uncertainty as a reference behavior, most grid tools are not readily prepared for performance monitoring of the AP as a whole, and specially detection of harmful events. In Figures 5-10, we propose a CL-EGA tool that is able to pinpoint deviations from a specified behavior corresponding to optimal control under uncertainty. Values along the y-axis represent *BG* levels obtained through a sensor whose readings are exposed to glycemic variability, different levels of calibration errors and time lags, insulin pump failures, and so on. The optimal policy describes the specified manner a glycemic regulator should behave; this specified behavior is reproduced in the x-axis. The specified *BG* dynamics incorporates the intrinsic (default) glycemic variability because of the glucose sensor errors and lag time for glucose absorption.

The proposed grid takes into account not only the difference between the measured *BG* values – of the implemented controller – and the expected *BG* values – of the specification – but also the clinical significance of this difference. Zone $\mathbf{F}^{-}$ represents severe hypoglycemia conditions. *BG* values falling

DOI: 10.1002/oca

in this zone result potentially life threatening – because of the risk of coma [29] – and are likely to require external assistance. Defective glucose counter regulation and hypoglycemia unawareness [30] episodes might lead to a vicious cycle of recurrent hypoglycemia and thus are used to characterize zone $\mathbf{E}^-$. Mild hypoglycemic episodes frequently precede severe hypoglycemia and are grouped in zone $\mathbf{C}^-$. Zones $\mathbf{A}^-$ and $\mathbf{A}^+$ both describe an optimally performing glucose regulator, that is, when the AP implementation closely follows the specification. Zone $\mathbf{B}^+$ denotes a state of increasing risk of progressing to diabetes but likely to revert to normal values. Zone $\mathbf{C}^+$ is classified as poorly controlled diabetes. Above these levels, for example, zone $\mathbf{D}^+$, demand aggressive control actions aiming at lowering *BG* to avoid any risk of a stroke. Strokes are the third leading cause of death and disability in the developed world, which make mandatory to maintain glycemic levels away from zone $\mathbf{E}^+$. For extremely high plasma glucose levels, the risk of diabetic coma is exceedingly high, requiring external assistance, and such a level is classified as grade zone $\mathbf{F}^+$. Note that those zones that describe hyperglycemic and hypoglycemic risk, that is, **E** and **F**, do not depend on meal intake and are limited by constant thresholds. A more detailed description of every grid zone is given in the Appendix.

3.3. Simulation of faulty scenarios

The specified behavior represents the optimally controlled glucose dynamics that the closed loop is expected to generate. This represents an optimal control policy acting over a regularly working AP (sensor, pump, control algorithm, and glucose metabolism). This closed-loop dynamics is represented by $\mathcal{GP}^*$ in (16). The specified dynamics $\mathcal{GP}^*$ – controlled by the optimal control policy for insulin infusion – is performed using variability $\sigma = 0.10$, calibration error $\xi = 2\%$, and time lag $\tau = 5$ min. The 24-h glucose profile is built using a sampling time of 6 min with the multiple meal consumption pattern shown in Table I.

A $\mathcal{GP}^g$ model describes the current behavior of the glycemic control loop (considering the effects of measurement errors, infusion pump malfunctioning, algorithm ill-tuning, and glycemic variability). This allows to evaluate the effect of a deviation from the specified behavior, which may be revealed by the CL-EGA tool. This Gaussian model describes an implementation of the glycemic regulator by a suboptimal control policy and possibly underperforming because of ill-functioning of the glucose sensors and the infusion pump. Suboptimal control strategies of the regulatory system are exemplified here by a proportional–integral–derivative (PID) algorithm and an integrated fuzzy-PID scheme. Even if different studies have shown a comparable performance between the control system and the routine clinical treatment [31, 32] using PID-type controllers, the ones presented in this work exhibit difficulties for efficiently controlling glycemic variability and are used here as representative examples of performance degradation. The intention is basically to display how an underperforming controller may be pinpointed by the monitoring tool, but a performance comparison between different glycemic regulators is out of the scope of the present work.

Low performance obtained through suboptimal control is (see *y*-axis on the grid) compared with the specified optimal behavior obtained using the LSMDP algorithm (see *x*-axis on the grid). This allows us to evaluate the overall efficiency of the AP when the control loop is affected by sensor errors as well as time lags compounded with the effect of intrinsic patient variability. The PID algorithm in (17) for the glucose–insulin model is fully described in Farmer *et al.* [33]. Here, u_b is the basal insulin, k_C is the proportional action, τ_I is the integral time, and τ_D is the derivative time. Parameters for the PID algorithm are given in Table III. To maintain basal conditions, the set point is set to the basal glucose concentration G_r.

$$I_{PID}(t) = u_b + k_C \left[\left(G(t) - G_r \right) + \frac{1}{\tau_I} \int_0^t \left(G(t) - G_r \right) dt + \tau_D \frac{d\left(G(t) - G_r \right)}{dt} \right] \tag{17}$$

The design of the expert fuzzy-PID controller is fully described in the work of Susanto-Lee *et al.* [34]. The fuzzy algorithm is based on a Mamdani-type fuzzy inference system, with one input variable (*BG* level) and one output variable (insulin delivery rate). The defuzzification for the output is calculated using the centroid method, which essentially determines the center of mass of the set of fuzzy outputs. For the proportional action, a sliding scale has a continuous blood glucose level partitioned into zones with a linear increment of the insulin rate according to the *BG* level. If the

DOI: 10.1002/oca

Table III. Proportional–integral–derivative parameters.

Parameter	Value
G_r	110 [mg dl^{-1}]
u_b	16.667 [mU min^{-1}]
k_C	12 [mU min^{-1} mg dl^{-1}]
τ_D	40 [min]
τ_I	3300 [min]

bolus relieved by the sliding scale does not provide enough insulin to lower the *BG* level, an integral actuator provides the extra increment needed to the insulin infusion rate. The derivative action is performed through a least square regression technique that provides a mechanism to boost the insulin delivery during a rapid increase in the *BG* level.

Figure 4 vividly highlights how suboptimal control policies affect the glycemic dynamics. More specifically, at the 12th hour, the parameter σ is significantly increased (from $\sigma = 0.10$ to 0.50) to simulate a sudden rise in glycemic variability. By contrasting the optimal and suboptimal controlled dynamics, the performance loss is readily revealed.

3.3.1. Closed-loop controller. The results of ill-tuning in the PID control strategy are portrayed in Figure 5. As the PID gain parameter k_C increases, so does the value of the output action for a given

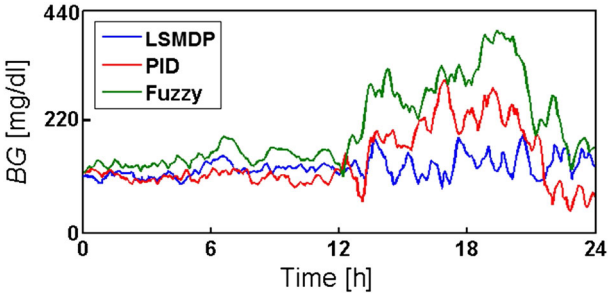

Figure 4. Effect of increasing glycemic variability using a noise scale $\sigma = 0.50$. LSMDP, linearly solvable Markov decision process.

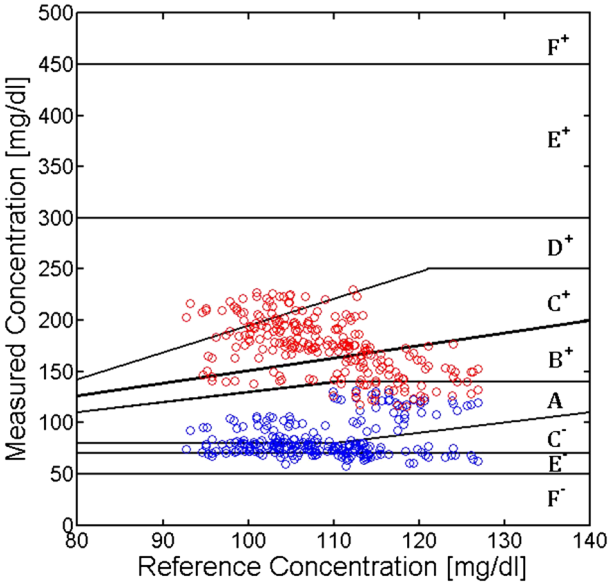

Figure 5. Performance degradation due to an ill-tuning of proportional action in the proportional–integral–derivative controller ($k_C = 20$ blue circles, $k_C = 5$ red circles).

DOI: 10.1002/oca

change in the error signal ($G(t)$-$G(r)$) in (17). However, excessive gain may lead the system to instability. In order to make the glycemic regulator capable of dealing with this scenario, the proportional parameter is increased from $k_C = 12$ to $k_C = 20$. A larger proportional gain causes glycemic levels to fall near the hypoglycemic limit, and this can be observed by the large number of blue circles falling below zone **A**. Later on, the proportional gain is reduced to $k_C = 5$, which makes the insulin infusion strategy much less effective giving rise to an increase in the *BG* levels (see red circles in Figure 5).

The same scenario was put into practice for the fuzzy-PID scheme in Figure 6, where the proportional gain k_P was first raised from 12 to 20 and later reduced to 5. Performance degradation in the glucose regulator is clearly revealed in the performance grid. Because of an excessive control action in the fuzzy-PID loop, the chances of a hypoglycemic event increase significantly (see blue circles). Similarly, a reduction of the proportional parameter k_P gives a severe performance degradation and *BG* reading shift towards hyperglycemic regions (red circles in Figure 6).

3.3.2. Glycemic variability. The capability of each controller to mitigate high levels of glycemic variability is at this point assessed. The value of the Ito's parameter of the glucose-insulin model is set to $\sigma = 0.50$, which gives rise to a much larger glycemic variability. A considerable reduction of the number of circles within zone **A** for all of the control techniques was observed, as it is shown in Figure 7. No dangerous *BG* circles can be observed when implementing the LSMDP algorithm (zones $\mathbf{E}^-$ and $\mathbf{F}^-$), despite enlargement of the spread of glucose levels (blue circles). For the PID strategy (red circles), *BG* values are evenly spread below and above the euglycemic zone, but some circles may indicate harmful situations for patients, as those falling in zone $\mathbf{E}^-$ and $\mathbf{F}^-$. Measured glucose concentration obtained by implementing the fuzzy-PID scheme indicates some difficulty to avoid hyperglycemia conditions, and several circles actually fall in zone $\mathbf{D}^+$ (green circles). The algorithm is not capable of calculating the right insulin bolus needed to mitigate such a level of glycemic variability in the patient. It is noteworthy how the performance of the fuzzy controller quickly degrades because of an increase in glycemic variability. However, in spite of this high variability, the robustness of LSMDP controller ensures that glucose levels fall inside the euglycemic range.

3.3.3. Continuous glucose monitoring system. The magnitude of the *BG-IG* lag may be no more than 5 min in optimal conditions, but after prolonged implantation, the sensor surface becomes increasingly fouled with fibrotic substance. As a result, the time lag progressively increases. In

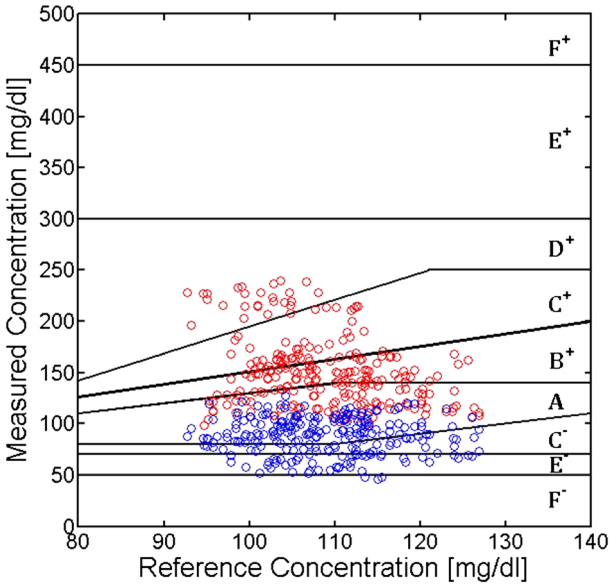

Figure 6. Fuzzy-PID ill-tuning as a result of varying the proportional gain ($k_C = 20$ blue circles, $k_C = 5$ red circles).

DOI: 10.1002/oca

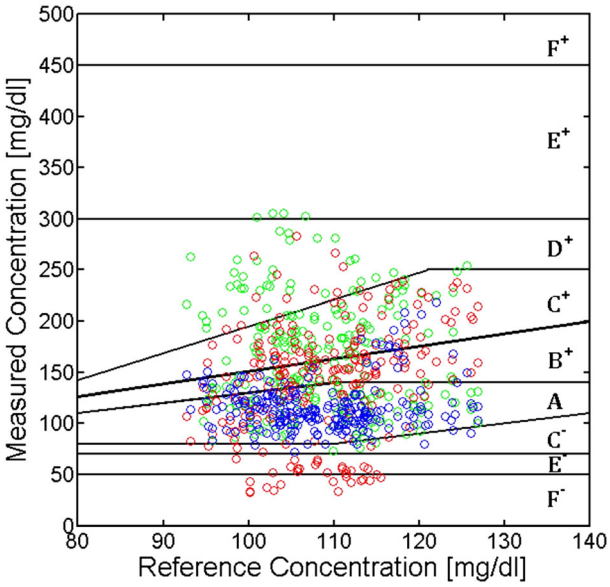

Figure 7. Assessment of the effect of high glycemic variability using a noise scale $\sigma = 0.50$ (linearly solvable Markov decision process blue, PID red, fuzzy-PID green).

Figure 1, truthful plasma glucose levels and their respective estimated sensor responses were depicted for different magnitudes of the calibration errors and time lags. The effect on the AP performance of a significant miscalibration of the glucose sensor ($\xi = 10\%$) as well as a large time lag ($\tau = 20$ min) are depicted in Figures 8 and 9, respectively. It is worth reminding that as the control algorithm responds to the outcome of the glucose sensor, it actually acts over inaccurate readings rather than true plasma levels. The same calibration errors and time lags do not give rise to a considerable performance degradation when implementing the LSMDP algorithm and most *BG* circles are in zones **A** and **B**. In turn, sensor readings affected by calibration errors and high time lags severely reduce the performance of a PID controller. As a result, the *BG* levels mostly fall in the

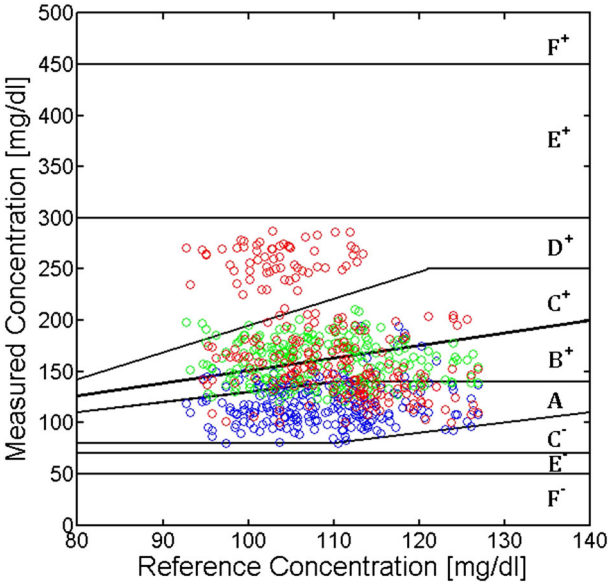

Figure 8. Performance loss due to glucose sensor miscalibration using $\xi = 10\%$ (linearly solvable Markov decision process blue, PID red, fuzzy-PID green).

DOI: 10.1002/oca

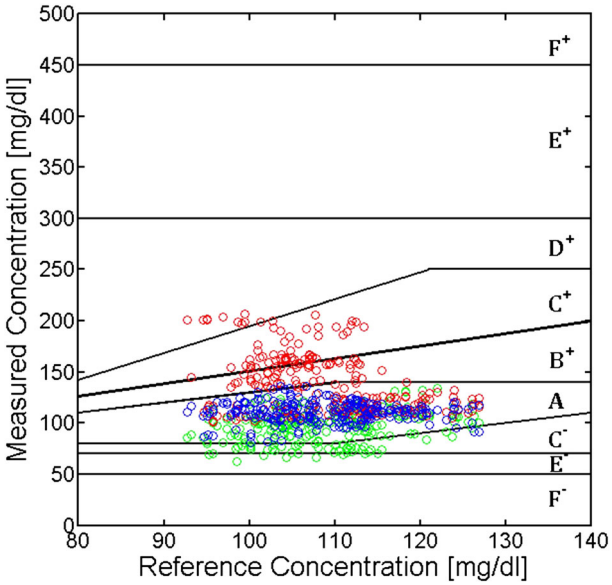

Figure 9. Performance loss due to a large time lag using $\tau = 20$ min (linearly solvable Markov decision process blue, PID red, fuzzy-PID green).

hyperglycemic range in an upper zone $\mathbf{D}^+$ of the grid. For the fuzzy controller, significantly higher calibration errors and time lags do not give rise to important deviations from accurate zones. A few *BG* circles in zone $\mathbf{D}^+$ and many more in zone $\mathbf{C}^+$ indicate poorly controlled diabetes with symptoms of hyperglycemia when there exist sensor errors and increased time lags. Hypoglycemia appears (green circles $\mathbf{C}^-$ and $\mathbf{E}^-$), when the sensor experiments an excessive time lag.

3.3.4. Continuous subcutaneous insulin infusion. Many commercially available insulin pumps deliver insulin continuously and subcutaneously. There exist local and systemic complications in controlled insulin delivery that include improper dosing because of electronic failures, catheter obstructions, battery depletion, and infections [35]. The simulation presented in Figure 10 tests a

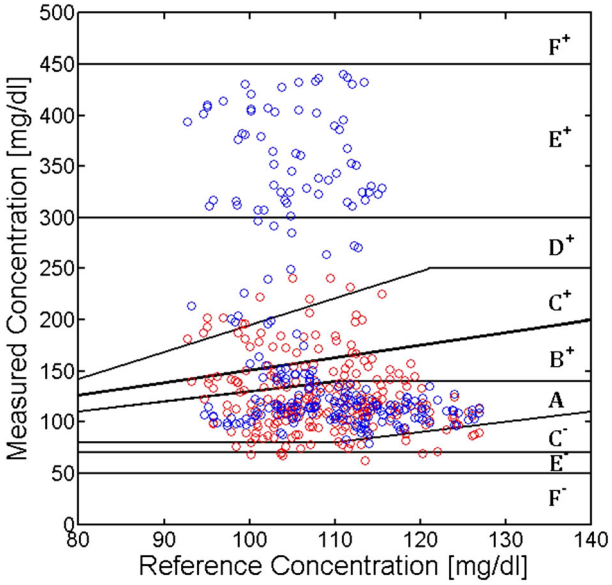

Figure 10. Performance degradation due to a catheter blockage (blue circles) and a coarse actuator (red circles) in an implantable insulin pump.

DOI: 10.1002/oca

likely scenario in which a catheter blockage occurs during the use of an insulin pump. Because of this blockage, insulin dose administrated via the pump suffers a 20% reduction with respect to the level prescribed by the optimal bolus calculated by the LSMDP algorithm. This reduces the effect of the bolus and leads to poorly controlled glycemic variability, which is portrayed by the amount of circles in the upper zones $\mathbf{D}^+$ and $\mathbf{E}^+$.

To simulate chaotic behavior of a coarser actuator (insulin pump), let us assume the actual infusion rate is given by $u(t) = u^* + \beta d\omega$, where u^* is the optimal action calculated using the LSMDP at time t, $d\omega$ is the differential of Brownian random noise, and $\beta = 3$ is the standard deviation of the added noise. Hence, the AP can calculate the optimal bolus to be administrated but is unable to set the right amount of insulin that is in fact delivered to the patient. Note that a Brownian noise is added to the optimal insulin action u^* to simulate the actuator fault in a realistic way. Because there is a zero mean noise, the path still converges to the vicinity of the target $BG = 110\,\mathrm{mg\,dl^{-1}}$, although there exists increased variability in the glucose dynamics. The chaotic controller has a similar effect as the one obtained by increasing the noise parameter σ, but here, excessive glycemic variability is caused by the regulator itself and not by physiological changes in the patient dynamics.

4. DISCUSSION

This paper proposes a novel grid-based monitoring tool to assess the performance of a closed-loop glycemic regulator. The main challenge for performance monitoring of an AP is defining the desired behavior under the uncertain conditions the glycemic regulator should face. To this aim, the monitoring tool has been tested on a stochastic model of a diabetic patient to assess its efficacy to pinpoint performance degradation under faulty conditions in sensor, control and insulin infusion pumps, or abnormal blood glucose variability because of a physiological disorder. The proposed monitoring tool highlights a well-established and user-friendly method for patients and clinicians such as the EGA in order to readily assess performance degradation. As the grid is divided in different zones, the CL-EGA allows fast detection of likely harmful episodes at a glance. Taking into account that most methods focus on performance assessment of glucose sensors, our performance grid is a one-of-a-kind attempt in the design of advanced tools aiming at monitoring an implantable AP.

The specification of the optimal glycemic regulator is obtained using a class of Markov decision processes that are linearly solvable. It has been reported [15] that the framework of LSMDP can find an optimal policy faster than conventional learning algorithms. Conversely, it requires the knowledge of state transition probabilities in advance, which perhaps is the most demanding optimization task. A major advantage of using the proposed approach is that it provides an explicit optimal policy for insulin delivery, which allows building a specification of the desired glycemic dynamics over time. Because the transition probability between different states is learned using GPs, it makes room for comparing the glucose readings with its specification for an optimally performing AP. Furthermore, the proposed performance monitoring tool has the advantage of being fully portable on a mobile device; because the optimal specification is provided explicitly, online optimization is not longer required.

In this paper, a monitoring tool is proposed. Monitoring works online and uses sensors to provide raw data and detect anomalies from the system under study whenever the system performs away from its expected behavior. Further analysis of the obtained data should be made using a diagnostic tool in order to analyze fault causes and provide meaningful information about the operation of the different system components [35]. Because both the expected glycemic variability and the sensor-related errors have been characterized through Gaussian transition probabilities, every significant deviation of points displayed in the grid is indicative of performance degradation in the control loop, for example, glucose sensor, control algorithm, pump, or even abnormal glycemic variability. Accordingly, our focus is in the AP as a whole; hence, the monitoring tool is instrumental to pinpoint any deviation from optimal performance.

To enhance the monitoring task, we are currently working on an online metric to detect deviations fast from an optimally performing AP. Furthermore, we intend to integrate learning the control policy for optimal acting with the monitoring tool. That is, currently, the monitor detects a deviant

DOI: 10.1002/oca

behavior, but no course of action for correcting it is taken. To address this issue, our current research efforts pursue tightly the integrating fault diagnosis within the glycemic regulator design in the framework of autonomic systems so that the controller can self-monitor, self-diagnostic, and self-optimize its control policy in real time.

APPENDIX. CLOSED-LOOP ERROR-GRID ANALYSIS

In CL-EGA, the values along the *y*-axis represent *BG* levels obtained through a sensor whose readings are exposed to glycemic variability, different levels of calibration errors and time lags, insulin pump failures, and so on. The specified manner a glycemic system should behave is reproduced in the *x*-axis. Fasting and prandial conditions must be taken into account to assess glycemic variability. A fasting state is a condition associated with no caloric intake for at least 8 h, whereas a prandial state describes a 2-h post-glucose intake containing the equivalent of 75 g of anhydrous glucose dissolved in water.

Plasma glucose at $50\,\mathrm{mg\,dl^{-1}}$ is used as threshold for zone $\mathbf{F}^-$, which represents severe hypoglycemia. *BG* values falling in this zone result potentially life threatening – because of the risk of coma [29] – and may require external assistance. Hypoglycemia antecedents, with prior plasma glucose concentrations as high as $70\,\mathrm{mg\,dl^{-1}}$, causes defective glucose counter regulation and hypoglycemia unawareness [30]. These episodes might lead to a vicious cycle of recurrent hypoglycemia, which set boundaries for zone $\mathbf{E}^-$. In addition to glucose counter regulatory systems that are triggered at this plasma glucose concentration in a nondiabetic person, warning symptoms of hypoglycemia are critical to allow interventions to restore plasma glucose concentration towards normal values. Mild hypoglycemic episodes frequently precede severe hypoglycemia. Diabetic persons have reported typical symptoms of hypoglycemia but with a measured plasma glucose concentration $>70\,\mathrm{mg\,dl^{-1}}$. Hence, zone $\mathbf{C}^-$ reflects the fact that patients with chronically poor glycemic control can experience symptoms of hypoglycemia at plasma glucose levels $>70\,\mathrm{mg\,dl^{-1}}$ [36]. Because of the proximity between the hypoglycemic symptoms in zone $\mathbf{C}^-$ and the hypoglycemic episodes in zone $\mathbf{E}^-$, a predictable zone $\mathbf{D}^-$ has not been included in the grid.

When blood glucose decreases to a concentration of about $80\,\mathrm{mg\,dl^{-1}}$, inhibition of insulin secretion occurs, resulting in a decrease in the use of peripheral glucose [37]. This value was set as the lower limit for normality during a fasting condition. The upper bound of $110\,\mathrm{mg\,dl^{-1}}$ was recommended as a limit for normality by the American Diabetes Association [38] for fast period. Notice that the range between 80 and $110\,\mathrm{mg\,dl^{-1}}$ defines a rectangle for accurate functioning in fasting conditions (zone $\mathbf{A}^-$) when the monitored AP operates optimally. Euglycemia during prandial conditions is restricted to the range from 110 to $140\,\mathrm{mg\,dl^{-1}}$. This range defines optimal functioning for a postmeal state (zone $\mathbf{A}^+$). Because glycemic risk is continuous, extending below and above the limit of the accurate range is needed; hence, two lines connect zones $\mathbf{A}^-$ and $\mathbf{A}^+$. This also avoids discontinuities in the grid and forms an inclusive zone **A**. Postprandial plasma glucose readings below zone **A** are likely to result from insulin overdosing. This situation may lead to an accelerated downward trend when glucose appearance from the gut is completed; thus, it is not advisable to incorporate a benign zone $\mathbf{B}^-$.

Zone $\mathbf{B}^+$ was restricted according to the World Health Organization diagnostic criteria for diabetes and intermediate hyperglycemia, through impaired glucose tolerance ($BG \geq 126\,\mathrm{mg\,dl^{-1}}$) and impaired fasting glucose levels ($BG \geq 200\,\mathrm{mg\,dl^{-1}}$) [39]. Thus, zone $\mathbf{B}^+$ denotes a state of increased risk of progressing to diabetes but likely to revert to normal values, which is actually referred to as pre-diabetes [40]. Furthermore, this state describes an intermediate group of subjects whose glucose levels, although not meeting the criteria for diabetes, are nevertheless too high to be considered normal. Plasma glucose above $142\,\mathrm{mg\,dl^{-1}}$ increases the risk of endothelial cell complications and is used as the upper limit for zone $\mathbf{C}^+$ under fasting conditions. The degree of endothelial dysfunction after a meal intake ranges from 142 to $300\,\mathrm{mg\,dl^{-1}}$ [41]. However, plasma glucose level above $250\,\mathrm{mg\,dl^{-1}}$ is classified as poorly controlled diabetes and thus set the upper limit for zone $\mathbf{C}^+$. Above these levels, ketones appear, and some restoring action aimed to lower *BG* is required; this corresponds to zone $\mathbf{D}^+$. In diabetic ketoacidosis, the plasma glucose concentration is typically

DOI: 10.1002/oca

greater than 250 mg dl^{-1}. The two major precipitating factors in the development of ketoacidosis are inadequate insulin treatment and infection. Although stroke is the third leading cause of death and disability in the developed world, no current standards exist for tight glycemic control, although the American Stroke Association recommends glucose to be maintained at $<$300 mg dl^{-1} [42]. This *BG* value is used as the defining threshold for zone $\mathbf{E}^+$. When plasma glucose is above 450 mg dl^{-1}, the risk of diabetic coma is exceedingly high, requiring external assistance, and such a level is classified as grade zone $\mathbf{F}^+$. Note that those zones that describe hyperglycemic and hypoglycemic risk, that is, **E** and **F**, do not depend on meal intake and are limited by constant thresholds.

REFERENCES

1. Aye T, Block J, Buckingham B. Toward closing the loop: an update on insulin pumps and continuous glucose monitoring systems. *Endocrinology and Metabolism Clinics of North America* 2010; **39**(3):609–624.
2. Renard E. Monitoring glycemic control: the importance of self-monitoring of blood glucose. *The American Journal of Medicine* 2005; **118**(9):12–19.
3. Patek SD, Breton MD, Chen Y, Solomon C, Kovatchev B. Linear quadratic Gaussian-based closed-loop control of type 1 diabetes. *Journal of Diabetes Science and Technology* 2007; **1**(6):834–841.
4. Ali S, Padhi R. Optimal blood glucose regulation of diabetic patients using single network adaptive critics. *Optimal Control Applications and Methods* 2011; **32**(2):196–214.
5. Westgard JO, Carey RN, Wold S. Criteria for judging precision and accuracy in method development and evaluation. *Clinical Chemistry* 1974; **20**(7):825–833.
6. Russell SJ. Continuous glucose monitoring awaits its "killer app.". *Journal of Diabetes Science and Technology* 2008; **2**(3):490–494.
7. Ghosh S, Maka S. A constrained sub-optimal controller for glucose regulation in type 1 diabetes mellitus. *Optimal Control Applications and Methods* 2014; **35**(2):191–203.
8. Hovorka R, Nodale M, Haidar A, Wilinska ME. Assessing performance of closed-loop insulin delivery systems by continuous glucose monitoring: drawbacks and way forward. *Diabetes Technology and Therapeutics* 2013; **15**(1): 4–12.
9. Harris TJ. Assessment of control loop performance. *The Canadian Journal of Chemical Engineering* 1989; **67**(5): 856–861.
10. Owens CL, Doyle III FJ. Performance monitoring of diabetic patient systems. In: *Engineering in Medicine and Biology Society*, *2001. Proceedings of the 23rd Annual International Conference of the IEEE.*Vol 2. IEEE; 2001: 2047–2050.
11. Huang B. Bayesian methods for control loop monitoring and diagnosis. *Journal of Process Control* 2008; **18**(9): 829–838.
12. Qing-wei M, Zhen-fang Z, Ji-zhen L. A practical approach of online control performance monitoring. *Chemometrics and Intelligent Laboratory Systems* 2015; **142**:107–116.
13. Naidu DS, Fernando T, Fister KR. Optimal control in diabetes. *Optimal Control Applications and Methods* 2011; **32**(2):181–184.
14. Todorov E. Eigenfunction approximation methods for linearly-solvable optimal control problems. In: *Adaptive Dynamic Programming and Reinforcement Learning*, *2009. ADPRL'09. IEEE Symposium on*. IEEE; 2009:161–168.
15. Todorov E. Efficient computation of optimal actions. *Proceedings of the National Academy of Sciences of the United States of America* 2009; **106**(28):11478–11483.
16. Dvijotham K, Todorov E. Linearly solvable optimal control. *Reinforcement Learning and Approximate Dynamic Programming for Feedback Control* 2012;119–141.
17. Lehmann ED, Deutsch T. A physiological model of glucose-insulin interaction in type 1 diabetes mellitus. *Journal of Biomedical Engineering* 1992; **14**(3):235–242.
18. Acikgoz U, Diwekar UM. Blood glucose regulation with stochastic optimal control for insulin-dependent diabetic patients. *Chemical Engineering Science* 2010; **65**(3):1227–1236.
19. Bergman RN, Phillips LS, Cobelli C. Physiologic evaluation of factors controlling glucose tolerance in man: measurement of insulin sensitivity and beta-cell glucose sensitivity from the response to intravenous glucose. *Journal of Clinical Investigation* 1981; **68**(6):1456–1467.
20. Rasmussen CE, Williams CKI. *Gaussian Processes for Machine Learning*. MIT Press: Cambridge, MA, 2006.
21. Bousquet-Rouaud R, Castex F, Costalat G, *et al*. Factors involved in catheter obstruction during long-term peritoneal insulin infusion. *Diabetes Care* 1993; **16**(5):801–805.
22. Jaremko J, Rorstad O. Advances toward the implantable artificial pancreas for treatment of diabetes. *Diabetes Care* 1998; **21**(3):444–450.
23. Richardson B. Diagnostics and monitoring of power transformers. In: *Condition Monitoring of Large Machines and Power Transformers* (*Digest No*: *1997/086*), *IEE Colloquium on*. IET; 1997:6/1–6/2.
24. Clarke WL, Cox D, Gonder-Frederick LA, Carter W, Pohl SL. Evaluating clinical accuracy of systems for self-monitoring of blood glucose. *Diabetes Care* 1987; **10**(5):622–628.
25. Magni L, Raimondo DM, Dalla Man C, *et al*. Evaluating the efficacy of closed-loop glucose regulation via control-variability grid analysis. *Journal of Diabetes Science and Technology (Online)* 2008; **2**(4):630.

DOI: 10.1002/oca

26. Chassin LJ, Wilinska ME, Hovorka R. Grading system to assess clinical performance of closed-loop glucose control. *Diabetes Technology & Therapeutics* 2005; **7**(1):72–82.
27. Zanderigo F, Sparacino G, Kovatchev B, Cobelli C. Glucose prediction algorithms from continuous monitoring data: assessment of accuracy via continuous glucose error-grid analysis. *Journal of Diabetes Science and Technology (Online)* 2007; **1**(5):645.
28. Sivananthan S, Naumova V, Man CD, *et al*. Assessment of blood glucose predictors: the prediction-error grid analysis. *Diabetes Technology & Therapeutics* 2011; **13**(8):787–796.
29. Ben-Ami H, Nagachandran P, Mendelson A, Edoute Y. Drug-induced hypoglycemic coma in 102 diabetic patients. *Archives of Internal Medicine* 1999; **159**(3):281.
30. Defining and reporting hypoglycemia in diabetes a report from the American Diabetes Association Workgroup on Hypoglycemia. *Diabetes Care* 2005; **28**(5):1245–1249.
31. Chee F, Fernando TL, Savkin AV, Van Heeden V. Expert PID control system for blood glucose control in critically ill patients. *Information Technology in Biomedicine, IEEE Transactions on* 2003; **7**(4):419–425.
32. Chee F, Fernando T, van Heerden PV. Closed-loop glucose control in critically ill patients using continuous glucose monitoring system (CGMS) in real time. *Information Technology in Biomedicine, IEEE Transactions on* 2003; **7**(1): 43–53.
33. Farmer TG, Edgar TF, Peppas NA. Effectiveness of intravenous infusion algorithms for glucose control in diabetic patients using different simulation models. *Industrial and Engineering Chemistry Research* 2009; **48**(9):4402–4414.
34. Susanto-Lee R, Fernando T, Sreeram V. Simulation of fuzzy-modified expert PID algorithms for blood glucose control. In: *Control*, *Automation*, *Robotics and Vision*, *2008*. *ICARCV 2008*. *10th International Conference on*. IEEE; 2008:1583–1589.
35. Gin H, Renard E, Melki V, *et al*. Combined improvements in implantable pump technology and insulin stability allow safe and effective long term intraperitoneal insulin delivery in type 1 diabetic patients: the EVADIAC experience. *Diabetes and Metabolism* 2003; **29**(6):602–607.
36. Boyle PJ, Schwartz NS, Shah SD, Clutter WE, Cryer PE. Plasma glucose concentrations at the onset of hypoglycemic symptoms in patients with poorly controlled diabetes and in nondiabetics. *New England Journal of Medicine* 1988; **318**(23):1487–1492.
37. Lacherade J-C, Jacqueminet S, Preiser J-C. Challenges in glycemic control in perioperative and critically ill patients: an overview of hypoglycemia in the critically ill. *Journal of Diabetes Science and Technology* 2009; **3**(6):1242.
38. Association AD. Diagnosis and classification of diabetes mellitus. *Diabetes Care* 2013; **36**(Suppl 1):S67–S74.
39. Alberti KG, Zimmet PZ. Definition, diagnosis and classification of diabetes mellitus and its complications. Part 1: diagnosis and classification of diabetes mellitus. Geneva: World Health Organisation. *Diabetic Medicine* 2004; **15**(7):539–553.
40. Gavin JR, Alberti KG, Davidson MB, *et al*. Report of the expert committee on the diagnosis and classification of diabetes mellitus. *Diabetes Care* 2003; **26**(Suppl 1):S5–S20.
41. Clement S, Braithwaite SS, Magee MF, *et al*. Management of diabetes and hyperglycemia in hospitals. *Diabetes Care* 2004; **27**(2):553–591.
42. Bloomgarden ZT. Inpatient diabetes control: rationale. *Diabetes Care* 2004; **27**(8):2074–2080.

DOI: 10.1002/oca