A model free approach to controlling blood glucose

L. Santoso I.M.Y. Mareels
Department of Electrical and Electronic Engineering
The University of Melbourne
Parkville, VIC 3010, Australia
l.santoso@ee.mu.OZ.AU, i.mareels@ee.mu.OZ.AU

Abstract

We present a problem of controlling type 1 diabetes mellitus - a situation where plant is complex and dynamic, the measurements are sparse, and the data display erratic fluctuating behaviour. These characteristics make it very difficult to derive a model of the plant. We propose a model-free method of deriving a controller for this problem and discuss the limits of control performance under the constraints.

1 Introduction

Diabetes mellitus is an age-old disease which until now remains incurable. It is a chronic disease characterised by the body's inability to regulate metabolism of sugars. This disease poses as a serious public health problem, as its poor management leads to complications which are debilitating to the individual and costly to the society. It is a case where a complex plant is relatively controllable by simple decision rules, making it an interesting control problem. We discuss in this paper the diabetes mellitus management problem as a case example of constrained control, where only little information is available to manage a complex and poorly defined situation. We present the derivation of a control algorithm directly from patient data.

Type 1 diabetes mellitus is a dysfunction of the glucose regulatory system, characterised by the absence of the hormone *insulin*. It is treated by subcutaneous injection of insulin. Currently, the typical management regime is multiple daily injection, prior to each meal and at bedtime. This treatment aims towards restoration of normal glucose levels, as in non-diabetic people, i.e. between 3.5 mmol/l and 7 mmol/l.

Avoidance of complications due to "bad" management is crucial, however achieving "good" management is not so trivial, as this control objective is not well defined. The glucose regulatory system is embedded in this very complex system we call the body. The state space is likely to be infinite, and only one variable, blood glucose (BG), is measured, sparsely. The state is not yet defined and is not observable. There is also wide inter-individual variations, bringing about the need for individualised rules. With these constraints in place, it is understood that the ideal aim of non-diabetic level of control is unachievable, leaving the realistic control goal as an open question.

We represent this system as follows:

$$\mathbf{x}_{k+1} = \mathcal{F}(\mathbf{x}_k, ins_k, F_k, E_k, \nu_k), \quad (1)$$

$$BG_k = \mathcal{G}(\mathbf{x}_k), \quad (2)$$

$$ins_k = \mathcal{H}(BG_k, \hat{E}_k, \hat{F}_k). \quad (3)$$

where $\mathbf{x}_k$ represents the state of the glucose metabolic system at time k, BG_k, ins_k, F_k, E_k denote BG measurement, insulin dose, food and exercise at time k, and ν_k is randomness which represent the unmeasured entities in the system. The "hats" represent estimated (effect) of food and exercise. BG is measured in mmol/l, ins in *Insulin Unit* (IU), F and E are qualitative measures. One of the questions addressed is the derivation of insulin rules $\mathcal{H}$, to control the complex system $\mathcal{F}$, where (at present) the control goal is unclear, with only meagre data points available. We take the approach of direct derivation, bypasssing the estimation of the "body" model $\hat{\mathcal{F}}$, for which we have insufficient information.

These issues are presented in this paper in the following manner. Firstly we define a control objective, based on the presently used clinical measure. We then discuss the framework of the insulin rules and how to customise it to an individual. To follow is the implementation of this algorithm by stochastic approximation methods.

2 Control objective

Even though BG is the only quantitative measurement taken, it is not consistently used to ascertain control performance. Clinicians in general accept that individual BG measurements are not good measure of a person's level of glucose control. Moreover, different practitioners often have different concerns and foci in their methods, possibly depending on the variety of cases they have been exposed to, resulting in lack of uniformity in opinions.

The general direction is towards the unattainable ideal situation, i.e. non-diabetic level of control. However, how far short of this ideal goal is considered acceptable, is an open question. It is also highly likely that the "best" level of performance is individual dependent.

A globally accepted goal is that from public health point of view, i.e. the reduction and prevention of late onset complications [6], which has been linked to level of $\mathrm{HbA_{1c}} \leq 7\%$ [1][1]. This benchmark has been used in clinical management as a goal.

In a non-diabetic scenario, blood glucose level would be maintained around the "normal" range. We employ as a criterion, an index which penalises blood glucose deviation from the normal region known as the *M-value* [3]:

$$M = 10\left|\log_{10}\frac{BG}{4.4mmol/l}\right|^3. \tag{4}$$

In keeping with clinical practice, we use a measure which, as $\mathrm{HbA_{1c}}$, is an indicator of BG average, the *Mean Amplitude Glucose Excursion* (MAGE) [3],

$$MAGE = \frac{1}{n}\sum_{k=1}^{n} 10\left|\log_{10}\frac{BG}{4.4mmol/l}\right|^3. \tag{5}$$

We relate MAGE to $\mathrm{HbA_{1c}}$ as follows. Let us denote an empirical mean of φ as $<\varphi>$,

$$<\varphi(x)> = \frac{1}{N}\sum_{k=1}^{N}\varphi_k(x), \tag{6}$$

$$\approx \int_0^{\infty}\varphi(x)p(x)dx, \tag{7}$$

where $p(x)$ is the probability density of x. Since $MAGE =< M(BG) > \approx E[M(BG)]$, and M is a real convex function [2],

$$MAGE(BG) =< M(BG) > \tag{8}$$

$$\geq M(< BG_k >), \tag{9}$$

from which it follows that

$$< BG > \in [\min M^{-1}(MAGE(BG)), \max M^{-1}(MAGE(BG))]. \tag{10}$$

$\mathrm{HbA_{1c}}$ has been correlated to a function of BG average over five weeks [5],

$$\mathrm{HbA_{1c}} = 2.07\times 10^{-2}\left(\frac{1}{N}\sum_{k=1}^{N}\frac{BG_k}{1\ mmol/l}\right)^{0.596} \tag{11}$$

[1] glycated haemoglobine, red blood cells to which glucose molecules are attached, used as a measure of long term blood glucose profile.

where N covers a five week period. Thus, in accordance to the clinical goal of $\mathrm{HbA_{1c}} \leq 7\%$,

$$2.07 \times 10^{-2}\left(\frac{< BG >}{1\ mmol/l}\right)^{0.596} \leq 7 \times 10^{-2} \tag{12}$$

$$< BG > \leq \left(\frac{7}{2.07}\right)^{\frac{1}{0.596}} mmol/l \tag{13}$$

$$< BG > \leq 7.72\ mmol/l. \tag{14}$$

This figure corresponds to $MAGE < 0.2$.

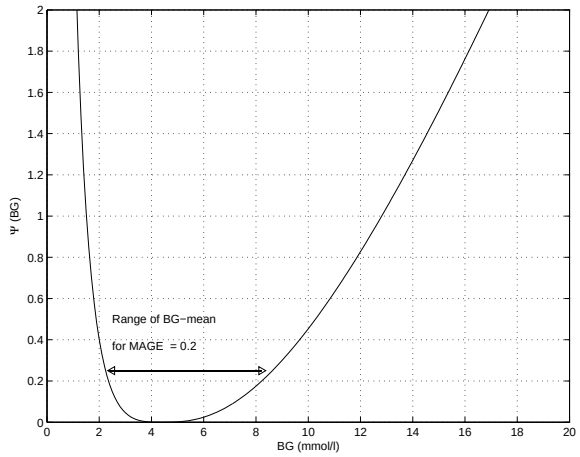

Figure 1: The criterion M-value function, showing the range of $< BG >$ for $MAGE \leq 0.2$.

3 Insulin rules

The relative success of clinical management of diabetes mellitus indicate that this task of controlling a complex system doable given the meagre amount of available data, using very simple control rules. We adopt this approach, and the rule framework used in clinical setting. We find that these rules, obtained from patient data, can be approximated by a piece-wise linear function.

We customise the rule to the individual and search for the best rule for that individual by parameterising it (Figure 2), as follows:

$$ins = Q[(ins_{min} + \frac{ins_{max} - ins_{min}}{BG_{max} - BG_{min}}.BG)(1 + \lambda.(F + E)]. \tag{15}$$

where the parameter corresponds to the break-points of the piece-wise linear curve

$$\boldsymbol{\theta} = (BG_{min}, BG_{max}, ins_{min}, ins_{max}) \tag{16}$$

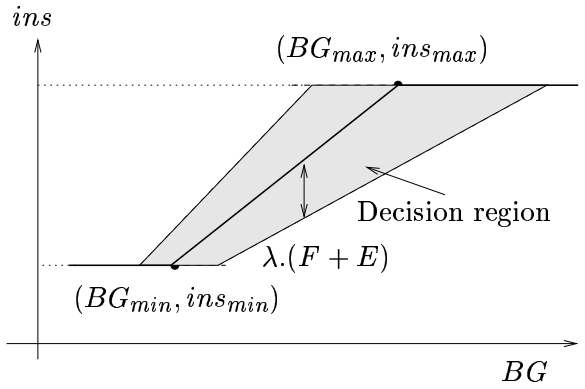

Figure 2: Parameterising the insulin rule by the break-points.

F, E denotes scaled effect of food and exercise, λ is a sensitivity parameter for food and exercise, and

$$Q(\cdot) = \begin{cases} ins_{min}, & ins \leq ins_{min}, \\ ins_{max}, & ins \geq ins_{max}, \\ round(\cdot), & otherwise. \end{cases} \tag{17}$$

A default value of $\lambda = 0.1$ is used in clinical management. Good performance is indicated by achieving low cost, as defined by the criterion function MAGE.

In this quest of the optimal rule parameters, there are obvious questions that must be settled; whether performance can be improved, and if so, how and by how much.

We address these questions by simulating a simple case of

$$BG_{k+1} = \mathcal{F}(BG_k, ins_k, \nu_k) \tag{18}$$

$$ins_k = \mathcal{H}_{\boldsymbol{\theta}}(BG_k) \tag{19}$$

where the diabetic plant is a function of immediate past BG and insulin only, ν_k represents randomness, and the insulin rule is only a function of glucose,

$$ins = Q[(ins_{min} + \frac{ins_{max} - ins_{min}}{BG_{max} - BG_{min}} BG)]. \tag{20}$$

The parameter $\boldsymbol{\theta}$ is as defined in Equation 16.

We want to find the parameter or set of parameters $\boldsymbol{\theta}^*$ for the best performing rule $\mathcal{H}_{\boldsymbol{\theta}}$ by direct adaptive control, as illustrated in Figure 3. The diabetic "subjects" are Markovian artificial patients, represented by Equation 18, and the rule is as in Equation 19. The rule performance is measured by the M-value function, whose output is taken by the tuner to update $\boldsymbol{\theta}$ towards $\boldsymbol{\theta}^*$.

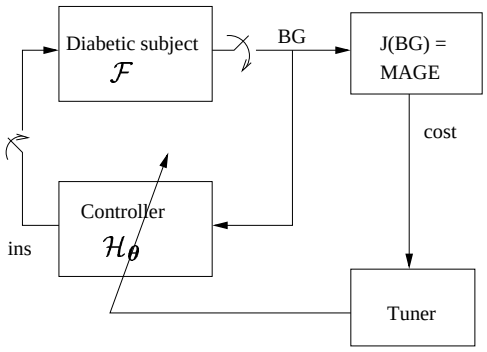

Figure 3: Direct adaptive control set-up to find the optimal rule parameter(s).

The known cost function $J(BG)$ is a function of blood glucose level. To find the optimal parameter, we need a cost function with respect of the parameter, $J(\boldsymbol{\theta})$, about which we have no direct information.

An intuitive idea of the minima locations would suggests a set of initial parameters for the adaptation algorithm. A knowledge of the number of minima region(s) would give us information on the feasibility of a gradient descent algorithm to find the global minimum.

We approximate the expected cost $E[J(\boldsymbol{\theta})]$ for a range of parameters $\boldsymbol{\theta}$ by the empirical mean $\frac{1}{N}\sum_{k=1}^{N} J_k(\boldsymbol{\theta})$ along a typical evolution of the controlled decision Markov Chain. This empirical mean is the MAGE function itself. With $N \approx 300$ a good approximation to $E[J(\boldsymbol{\theta})]$ is attained. The numerical estimation of the expected value is conducted over 100 runs, each of 300 iterations. This is equivalent to calculating MAGE with $K = 300$ (i.e. over ten weeks), and averaging 100 MAGE values.

The cost surface $E[J(\boldsymbol{\theta})]$ is a four dimensional hypersurface. For visualisation purposes, we show projections of the whole cost surface, fixing (ins_{min}, ins_{max}). The expected values were computed for the following pairs of (ins_{min}, ins_{max}): $ins_{min} \in \{0, 3, 5\}, ins_{max} \in \{10, 15, 20\}$. The experiment is conducted on Markovian subjects *I, H* and *O*, which represent the behaviour of ideal, hyperglycaemic and oscillatory subjects respectively. The latter two respective cases denote patients with tendencies of high BG levels and erratic BG fluctuations.

Figure 4 shows the cost surfaces for subjects *I, H* and *O*, for insulin boundaries $(ins_{min}, ins_{max}) = (0, 10)$. We can see that not only the shape is preserved, but the region of minimum cost is also located in the same area ($0 \leq BG_{min} \lessapprox 2\ mmol/l, 0 < BG_{max} \lessapprox 5\ mmol/l$). The minimum expected costs are higher for subject H and O, as anticipated.

These simulations thus far suggest that the cost surface calculate using the $MAGE$ function

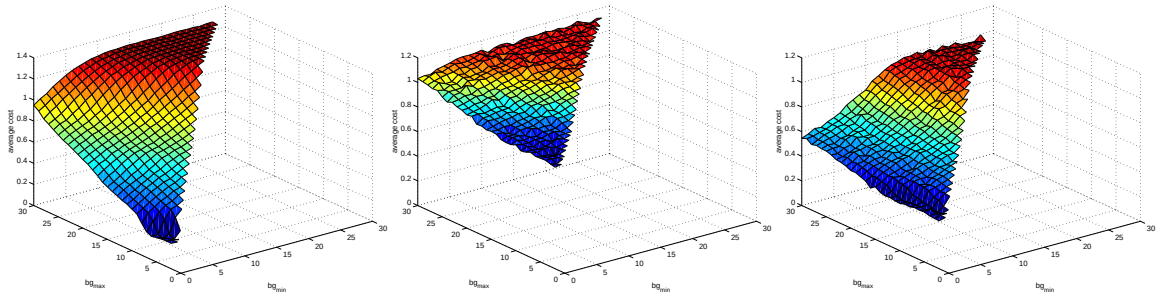

(a) Ideal subject (b) Hyperglycaemic subject (c) Oscillatory subject

Figure 4: Cost surfaces pertaining to three subject characteristics, for rules with $(ins_{min}, ins_{max}) = (0, 10)$ and $0 \leq BG_{min} \leq BG_{max} \leq 30$.

have the following properties:

- monotonicity,
- unique minimum region.

which lend themselves to safe implementation of a gradient descent method to find the optimal rule parameters. These properties are **independent** of the "person". Different subjects of different levels of "innate" controllability naturally yield different levels of optimum performance. However, the parameter area yielding for different subjects, under different rules, is not very sensitive to the characteristics of the individual.

4 Stochastic approximation

A stochastic approximation of the steepest descent algorithm is implemented to locate the optimal parameter, or neighbourhood of optimal parameters.The parameter vectors $\boldsymbol{\theta}$ are updated in the direction of negative gradient to the cost function J,

$$\boldsymbol{\theta}_{k+1} \;=\; \boldsymbol{\theta}_k \;-\; \alpha_k . \nabla J(\boldsymbol{\theta}_k). \tag{21}$$

We need to estimate the gradient of this unknown cost function $J(\boldsymbol{\theta})$, to determine the direction

we should optimise. We estimate the gradient of the cost function by random perturbation [4]

$$\nabla J_k \;=\; \frac{J(\boldsymbol{\theta}_k \;+\; \gamma_k.\Delta) - J(\boldsymbol{\theta}_k \;-\; \gamma_k.\Delta)}{2\gamma_k} \;.\; \begin{bmatrix} \Delta_1^{-1} \\ \vdots \\ \Delta_n^{-1} \end{bmatrix} \tag{22}$$

where the perturbation vector $\boldsymbol{\Delta}$ is randomly generated, with elements of a Bernoulli ± 1 distribution.

Bearing in mind that the plant $\mathcal{F}$ is stochastic, we need to update in a direction negative gradient on average. We can approximate the effect of averaging by updating with small step sizes α_0. The magnitudes of adaptation step size α and perturbation size γ_0 can initially be not so small (in the order of 1 mmol/l and 1 Insulin Unit respectively), and decrease monotonically

$$\alpha_{k+1} = \alpha_\infty + \frac{\alpha_0}{(A+k+1)^\phi}, \tag{23}$$

$$\gamma_{k+1} = \gamma_\infty + \frac{\gamma_0}{(k+1)^\psi}, \tag{24}$$

$$0 < \phi, \psi < 1. \tag{25}$$

Initially, it is relatively safe to perturb and update more boldly, but as the algorithm hopefully approaches the minimum cost region, the parameters should be more cautiously updated. The step and perturbation sizes are kept above a lower bound α_∞ and γ_∞, such that the algorithm never really terminates.

From the series of experiments we observe that in order to arrive in the optimal region and remain there, we need to pay attention to the step sizes and initial parameters. Safe values for initial parameters can be guessed from the cost surfaces we explored. We find that $BG_{min} = 0, BG_{max} = 10\ mmol/l$ (accepted border for normal-high BG levels), $ins_{min} \approx 0.5 ins_p$ and $ins_{max} = 2 ins_p$ are reasonable safe values to use.

The main caution to note, with regard to choices of initial parameters, is to avoid the high cost plateaux in $J(\boldsymbol{\theta})$. If the adaptation is started with an initial parameter $\boldsymbol{\theta}_0 = (0, 10, 15, 20)$, for example, it is likely to remain in that high cost plateau (see Figure 4 for reference). However, the above parameter corresponds to a rule which is utterly nonsensical. This rule advices patients to take no insulin unless their BG level exceeds 15 mmol/l. No diabetic person would survive very long under this insulin rule!

It seems that most sensible rules should sit in the region of non-zero gradient, hence enabling the algorithm to move towards parameters of lower costs.

We find that even if the unsensible initial parameter is applied, it is still possible to escape that area of maximum cost, by using relatively large initial step size. Small step size is typically wise, for a conservative search and to approximate averaging. However, if the initial parameter sits in a maximum plateau, larger steps are necessary to jolt the adaptation out of that area. Thus it is appropriate to initiate the algorithm with step size in the order of one unit (both in glucose and insulin), but then the step size should decrease as the algorithm proceeds. As shown in Equation 23, the step size decreases as $k^{-\phi}$, where k is the iteration index and $0 < \phi < 1$.

We need to bear in mind as well that the real system is likely to be time variant. The optimal parameter may shift, so it is a good idea to keep the algorithm "searching". Thus we need to place a lower bound to the step size α, such that $\alpha \nrightarrow 0$ as $k \to \infty$. This strategy has a disadvantage, that if the optimal parameter does not actually change, then the algorithm would not reach it. If it reaches it at a point in time, it would move away again. However, this is not so much of an issue, if α_{min} is kept small. As we see in the surface plots, that the family of optimal parameters cover a reasonably sized neighbourhood. Small movements around that neighbourhood is unlikely to cause movements towards significantly higher cost regions. The benefit of keeping the algorithm alive is worth this small cost.

5 Discussion and concluding remarks

We have conceptually shown that it is feasible to control a complex system as glucose metabolism in a constrained situation, by means of direct control derivation. This approach is relatively simple, bypassing the need of constructing a model of the complex plant, for which we have insufficient data. This approach is also consistent with the clinical management scheme.

This situation pose lack of clarity in a few aspects, about the state of the system and the control goal itself. The state is not defined, and is likely to be infinite. In any case, there is inadequate information to observe the state. This direct approach allows us to derive the control without a detailed knowledge of the plant state.

However, it is crucial to formulate a definite control goal. We infer a control objective in terms of blood glucose level from the current clinical scheme benchmark, the $\mathrm{HbA_{1c}}$. The criterion we employ aim to minimise glucose deviations, taking into account both the mean and variance of blood glucose.

Thus far, HbA_{1c} has only been related to BG-mean. Taking this interpretation means that the current clinical management only aims to control BG-mean, and not paying much attention to the variance. This is not in line with the actual practice, where wide BG fluctuation is considered equally undesirable as a high mean. There is no strict lower bound to this goal. If $HbA_{1c} \leq 7\%$ were to be taken as a control goal, an optimiser would look for the lowest possible value, potentially causing glucose to plunge too low.

The $MAGE$ function has an advantage, that it penalises both low and high glycaemic levels. Widely fluctuating glycaemic levels give a high $MAGE$ score, even if the BG-mean sit in the range specified for that score. This makes the $MAGE$ function an appropriate penalty function with which to formulate a control goal. It reflects the populational measure, it allows continuous monitoring of control performance, and it penalises both mean and variance of BG levels, consistent with the day-to-day practice of BG management.

Thus far, our algorithm has been trialled only in simulation "subjects". Simulation experiments show that the algorithm is capable of finding the rules which would yield the "optimal" control performance, as measured by the MAGE cost function. Whether it would perform similarly when a person is involved, instead of a Markov model, remains to be seen.

There are remaining questions which can only be answered by clinical trial on diabetic subjects. Would the possibly more complex system, such as a real human body, affect the performance of the algorithm? Would it still find the "optimal" rules in reasonable time? It would be particularly interesting to observe the algorithm performance on a brittle diabetic. Would it actually make a difference to the person, and by how much?

There are remaining non technical issues if such a device is to be released in the market. Firstly, a framework for the legal liabilities need to be established. Secondly, there is a question of user acceptance. Thus far we have encountered two practically opposite camps, with regard to the usage of an advisory device, those who want it, and those who don't. We encounter diabetic people who are at lost with the uncontrollability of their glycaemia and people who are far away from medical attention. These people are very enthusiastic to accept developments which offer assistance, and a hope for a better glycaemic profile.

However, there are those who believe that diabetics should understand their own bodies. Using technologial aid may lead to dependency to such devices, and rob them of the empowerment of this understanding. These people typically have better control of their own BG level, and have

some sort of support group in their community.

References

[1] Diabetes Control and Complications Trial Research Group. Resource utilisation and cost of care in the diabetes control and complications trial. *Diabetes Care*, 18:1468–1478, 1995.

[2] J. Neveu. *Bases mathématiques du calcul des probabilités*. Masson et C, Paris, France, 1964.

[3] F.J. Service, G.D. Molnar, J.W. Rosevear, E. Ackerman, L.C. Gatewood, and W.F. Taylor. Mean amplitude of glycemic excursions, a measure of diabetic instability. *Diabetes*, 19:644–655, 1970.

[4] J.C. Spall. An overview of the simultaneous perturbation method for efficient optimization. *John Hopkins APL Technical Digest*, 19:482–492, 1998.

[5] P.A. Svendsen, T. Luntzen, U. Sœgaard, and J. Nemp. Glycosylated haemoglobine and steady state mean blood glucose concentration in type I (insulin dependent) diabetes mellitus. *Diabetologia*, 23:403–405, 1982.

[6] World Health Organisation (Europe) and International Diabetes Federation (Europe). Diabetes care and research in Europe: the St Vincent Declaration. *Diabetic Medicine*, 7:360, 1990.