

Continuous Performance Monitoring of a Closed-Loop Insulin Delivery System using Bayesian Surprise

Luis Ávila, Ernesto Martínez

INGAR (CONICET-UTN), Avellaneda 3657, Santa Fe S3002 GJC, Argentina
{avilaol, ecmarti}@santafe-conicet.gob.ar

Abstract. The development of an artificial pancreas for the treatment of insulin-dependent diabetes is a big challenge for control theory. Many closed-loop algorithms have been widely evaluated for their ability to recreate, as closely as possible, glucose and insulin profiles observed in healthy individuals. Nevertheless, an artificial pancreas system also involves a critical necessity for supervision of the control loop functioning and the early detection of anomalies as well as performance deterioration. This work presents a performance monitoring approach using Bayesian surprise to fast detect functional degradation and guarantee adequate control of blood glucose levels. Bayesian surprise is significantly affected by any deviation from desired operation in a controlled system, which allows its use for continuous glucose monitoring.

Keywords: Diabetes, optimal action selection, bayesian surprise, artificial pancreas, continuous performance monitoring.

1 Introduction

Type 1 diabetes mellitus (T1DM) is a disease that results in lifetime dependence on exogenous insulin due to the inability to naturally produce insulin in the pancreas. As the ability to produce insulin is completely lost, those with T1DM must receive regular insulin infusions to survive. The aim of glucose control using insulin infusions is to maintain *normoglycemia*, i.e., a blood glucose level (BGL) between 3.9 and 7.8 mmol/L. Untreated T1DM primarily causes elevated blood glucose levels and acidosis. Intensive therapies managed by multiple daily injections and monitoring of BGL, has had success, but remains an annoyance to the subject and often results in poor control. In order to prevent these complications and achieve a better quality of life for diabetic patients, effective regulation of blood glucose is essential. The ideal treatment for controlling BGLs would be the use of an artificial pancreas. Such fully automated system shown in Fig. 1 would consist on a glucose sensor to monitor the blood glucose continuously with sufficient reliability and precision, a controller to calculate the necessary insulin infusion rates by an appropriate feedback algorithm and an insulin infusion pump to supply the required amount of insulin into the blood.




Fig. 1. Closed-loop insulin delivery system

To achieve an optimally-controlled system, a clinically implantable artificial pancreas requires well functioning and coordination of all their components (Jaremko and Rorstad, 1998). However, certain problems arise from different angles. Firstly, substances with absorption spectra similar to glucose incur complications due to obstruction of the infusion catheter or sensor unstable signal output. The same way, changes during the material lifecycle or maladjustment of operation strategy may lead to the degradation of control performance. A stabilized design of the integrated artificial pancreas should guarantee well functioning beyond uncertain changes in the loop behavior. These drawbacks make clear the necessity for supervision of the control loop and the early detection of performance deterioration.

To this aim, optimal choice of actions is a fundamental problem to be addressed in order to characterize the behavior of an optimally-controlled system. The abstract setting for the latter can be framed as an agent (or controller) choosing actions over time, an uncertain dynamical system whose state is affected by those actions, and a performance criterion that the agent seeks to optimize (Todorov, 2009). The probabilistic controller has the power to reshape the system dynamics in any way it wishes. However, it pays a price for too much reshaping (Kárný and Guy, 2006). The key question for loop monitoring is how the “distance” from optimal reshaping can be characterized using realizations of a stochastic process bearing in mind optimal control of a nonlinear system.

Optimal actions are defined by the stochastic dynamics $p(x'|x, u)$ specifying the optimal transition probability from state x to state x' under control action u . As the control system implements a given controlled-system dynamics $g(x'|x, u)$ which deviates from optimal operation its performance necessarily degrades. The rationale behind $p(x'|x, u)$ is the optimal cost-to-go function $v(x)$, defined as the expected cumulative cost for starting at state x and acting optimally thereafter. It compresses all relevant information about the future and thus enables greedy computation of optimal actions. The value function $v(x)$ equals the minimum of the immediate cost plus the expected cost-to-go $E[v(x')]$ from the next state x'

$$v(x) = \min_u \{ \ell(x, u) + E_{x' \sim p(\cdot|x, u)} [v(x')] \} \quad (1)$$

where the subscript indicates that the expectation is taken with respect to the transition probability distribution $p(\cdot|x,u)$ induced by action u . Eq. 1 is fundamental to optimal control theory and is called the *Bellman* fundamental equation.

The cost of reshaping the dynamics of a system can be measured with reference to the *passive* dynamics characterizing the behavior of the system in the absence of controls. The latter behavior will usually be defined as a random walk in discrete domains and as a diffusion process in continuous domains. Note that the notion of passive dynamics is common in continuous domains but is rarely used in discrete domains. Let $\wp(x'|x)$ denote the passive system dynamics. The control system can influence this dynamics in any way it wishes. However, it pays a price for loop reshaping beyond what is strictly necessary for optimal control. The minimum immediate cost for optimal reshaping can be estimated as follows

$$\ell^*(x, u) = q(x) + \mathbf{E}_{x' \sim p(\cdot|x,u)} \log \left[\frac{p(\cdot|x,u)}{\wp(x'|x)} \right] = q(x) + KL(p(x'|x, u) \parallel \wp(x'|x)) \quad (2)$$

The state cost $q(x)$ is an arbitrary function encoding how (un)desirable different states are and KL is the Kullback–Leibler divergence that measures the difference between the optimally-controlled dynamics and the passive one. Eq. 2 can be written for any arbitrary shaping associated with $\wp(x'|x, u)$ which allows modeling the performance loss for too much reshaping the passive dynamics as follows

$$\ell^\wp(x, u) - \ell^*(x, u) \cong KL(\wp(x'|x, u) \parallel p(x'|x, u)) \quad (3)$$

The KL divergence is also known as the *relative* entropy or the Kullback–Leibler (KL) entropy, which may be understood as a measure of the difficulty of discriminating between two distributions. Entropy is a quantity defined for any probability distribution with properties that agree with the intuitive notion of information content and in stochastic optimal control can be directly related to the approximate solution of the Fokker-Planck-Kolmogorov equation (Plastino, et al., 1997; Günel, 2010). Also, Majda et al., (2002) have recently demonstrated that the relative entropy provides a significant measure of the information content of a prediction ensemble. More specifically, entropy and *predictability* are intertwined concepts relevant for performance loop monitoring (Ghraizi et al., 2009).

The distance between two distributions in Eq. 3 is increased by any deviation from optimal operation in process systems which makes it ideal for supervising Real-Time Optimization (RTO) and dynamic optimization schemes.

The earlier bayesian definition for surprise proposed in Baldi et al. (2010) quantifies how observing new data affects the internal beliefs that an observer may have over a set of hypotheses or models of the world. Surprise is measured using the distance between the posterior and prior distributions, based on the Kullback-Leibler divergence. Hence, *Bayesian Surprise* is perturbed by any disappearance as well as appearance of stimulus from operation in process systems, which makes it ideal for supervising RTO and dynamic optimization schemes. In these types of control tasks, performance cannot be characterized aptly enough by classical concepts such as *min-*

imum variance (Owens and Doyle, 2001) or the *predictability of error residuals* (Ghraizi et al., 2009).

2 Methodology

The Bergman's two-compartment minimal model parameterized as described in Acikgoz and Diwekar (2010) is used to model the glucose-insulin dynamics in a simulated patient with proper addition of an Ito's stochastic process to capture the patient variability using the variance parameter σ . Significant variability of relevant parameters among patients and within a given patient during the course of the day or week affects the insulin and glucose dynamics. Commonly this dynamics reveals some type of organization or behavior dissimilar from uniform randomness. The ability to identify patterns of behavior allows the characterization of the realizations of a stochastic process control, anomaly detection, decision making and driving actions.

2.1 Bayesian Surprise

Surprise is strong in real-life environments when an abruptly change causes a reevaluation of beliefs about the nature of the system. Because surprise exists only in the presence of uncertainty its essence must be probabilistic, using distributions to capture subjective expectations or beliefs over the current space of hypotheses or models M . These beliefs are updated, as data is acquired, transforming prior belief distributions into posterior belief distributions. The fundamental effect of the data D on the observer is to change the prior distribution $P(M)$ into the posterior distribution $P(M|D)$ via Bayes theorem

$$P(M|D) = \frac{P(D|M)}{P(D)} P(M) \quad (4)$$

The new observation D carries no surprise if it leaves the observer's beliefs unaffected, that is, if the posterior is identical to the prior; conversely, D is surprising if the posterior distribution resulting from observing D significantly differs from the prior distribution. Surprise is measured using the distance between the posterior and prior distributions, based on the Kullback-Leibler divergence.

$$S(D, M) = KL(P(M), P(M|D)) = \int_M P(M) \log \frac{P(M)}{P(M|D)} dM \quad (5)$$

To compute surprise, when the prior and the posterior distributions have the same functional form, we need to calculate general terms of the form

$$F(P_1, P_2) = \int P_1 \log P_2 dx \quad (6)$$

where $P_1(x)$ is the prior and $P_2(x)$ is the posterior. Surprise is then given by

$$S = F(P_1, P_1) - F(P_1, P_2) \quad (7)$$

Consider a data set $D = \{d_1, \dots, d_n\}$ containing N points. When d_i is real and data D has unknown mean and unknown variance, we have a family $M(\mu, \sigma^2)$ of models with a prior

$$G_1 \Gamma_1 = G_1(\mu_1, \sigma^2/k_1) \Gamma_1(v_1, s_1) \quad (8)$$

product of a normal with a scaled Gamma distribution. Thus the prior and the posterior have four parameters (μ, k, v, s) satisfying (Gelman et al., 1995)

$$\begin{aligned} \mu_2 &= \frac{k_1}{k_1+N} \mu_1 + \frac{N}{k_1+N} \bar{m} \\ k_2 &= k_1 + N \\ v_2 &= v_1 + N \\ v_2 s_2^2 &= v_1 s_1^2 + (N-1) \bar{\sigma}^2 + \frac{k_1 N}{k_1+N} (\bar{m} - \mu_1)^2 \end{aligned} \quad (9)$$

where \bar{m} and $\bar{\sigma}^2$ are given by

$$\begin{aligned} \bar{m} &= \sum d_i / N \\ \bar{\sigma}^2 &= \sum (d_i - \bar{m})^2 / (N-1) \end{aligned} \quad (10)$$

The Bayesian surprise is subsequently computed as

$$\begin{aligned} S(D, M) &= \frac{1}{2} \log \frac{k_1}{k_1+N} + \frac{N}{2k_1} + \frac{k_1+N}{2} \left[\frac{N(\bar{m}-\mu_1)}{(k_1+N)s_1} \right]^2 + \log \frac{C_1}{C_2} - \\ &\left[\frac{N}{2} \psi \left(\frac{v_1}{2} \right) + \log \frac{2}{v_1 s_1^2} \right] + \frac{(N-1) \bar{\sigma}^2 + \frac{k_1 N}{k_1+N} (\bar{m}-\mu_1)^2}{2s_1^2} \end{aligned} \quad (11)$$

where ψ is the derivative of the logarithm of the Gamma function. For large values of N it yields to the approximation

$$S(D, M) \approx \frac{N}{2} \left(\frac{1}{k_1} + \frac{\bar{\sigma}^2}{s_1^2} + \log \frac{v_1 s_1^2}{2\bar{\sigma}^2} - \psi \left(\frac{v_1}{2} \right) + \frac{(\bar{m}-\mu_1)^2}{s_1^2} \right) \quad (12)$$

2.2 Closed-Loop Control Strategies

The performance of closed-loop insulin infusion algorithms is assessed implementing Bergman minimal model describing insulin and glucose dynamics within a type I diabetes patient. The well-known Bergman's two-compartment minimal model is used to model the glucose-insulin dynamics in a simulated patient with addition of an Ito's stochastic process to capture the patient variability. Table 1 shows the schedule for carbohydrate intake.

Table 1. Meal times for carbohydrate intake.

Times (min)	180	300	450	660	870	1020
Carbohydrate (g)	47	16	63	31	63	31

Gaussian Process Dynamic Programming (GPDP) is a generalization of Dynamic Programming (DP)/Value Iteration (VI) to continuous state and action spaces using fully probabilistic Gaussian Processes (GP) models (Deisenroth et al., 2009). GPDP describes value functions directly in function space by representing them using fully probabilistic GP models that allows accounting for uncertainty in simulation-based optimal control. An algorithm of the GPDP algorithm using the transition dynamics GP_f and Bayesian active learning was presented by De Paula et al. (2012).

Surprise theory can further be used in experiments where the prior is biased by an antecedent in such a manner that it is possible to represent a reference control policy as prior exposures to stimuli. Thereby, surprise can be helpful for quantitative comparison between different control schemes that may have different policies or priors. The optimal control policy obtained through the GPDP algorithm using $\sigma=0.10$ is used as reference to assess the behavior of PID, expert and fuzzy schemes with respect to maintaining normal BGLs during a meal. PID method is based on the work presented by Farmer et al. (2009) of a simple proportional plus integral plus derivative (PID) feedback control algorithm. The design of an Expert PID controller is based on the work developed in of Chee et al. (2003), upon a sliding scale method which prescribes insulin rates based on a lookup Table. The table has a continuous BGL partitioned into zones with linear increments of insulin rate, implementing a proportional action. If the current amount of insulin is not able to lower BGL, an integral control mechanism provides an increment of insulin through a variable which depends on a normalized weighted average of the latest BGLs. The derivative action is performed by a least squares regression technique to boost the insulin delivery during a rapid increase of BGL. The insulin infusion dose is calculated upon proportional, integral and derivative actions. The described expert algorithm is modified through the addition of fuzzy membership functions to obtain the proportional, integral and derivative actions (Susanto-Lee et al., 2008). The controller is based on Mamdani-type architecture with defuzzification of the output variable implemented by the centroid method. Proportional, integral and derivative actions are replaced by membership functions and the infusion dose is calculated as in the expert scheme.

3 Results

Fig. 2 depicts 1-day simulation (1440 min.) of the glucose controlled dynamics using a generic control policy π_G obtained from the GPDP algorithm for blood glucose. Patient response variability is simulated by different values of the Ito's variance parameter σ and the carbohydrate intake given by Table 1. It is worth noting how the performance degrades as the variance parameter σ is gradually increased. In this way,

model parameters are drastically changed so as to obtain a somehow sub-optimal control policy.

BGLs obtained for PID, Expert and Fuzzy control schemes are shown in Fig. 3. It is notably the difference between each controlled glucose dynamic. While the controller obtained using the GPDP algorithm presents rather smooth actions, remaining schemes behave mostly as on-off control systems in response to the Ito's process variability.

As the variance parameter σ increases the major variability causes glucose levels to fall outside of the normoglycemic range. This type of malfunctioning is notably captured by the Bayesian surprise as shown in Fig. 4. As previously stated, new glucose data carries no surprise if BGLs remain mainly constant leaving prior distributions unaffected, as in the case of variance $\sigma=0.10$. Conversely, as σ increases the arrival of surprising data modify prior/reference distributions causing larger surprise levels. Similar behaviors can be inferred from Fig. 5. The on-off response obtained from fuzzy and expert controllers is reflected by surprise levels. Note that PID scheme presents a stationary surprise value due to accumulation of constantly new surprising data.




Fig. 2. Blood glucose profiles for different values of variance parameter σ




Fig. 3. Blood glucose profiles for different control schemes using $\sigma=0.10$




Fig. 4. Bayesian surprise for blood glucose profiles for an increment in σ




Fig. 5. Bayesian surprise for blood glucose profiles using different control schemes




Fig. 6. Surprise for a change in variance parameter σ at min. 1440




Fig. 7. Surprise for a change of control scheme at min. 1440




Fig. 8. Change in diet and control policy for the fuzzy scheme

In Fig. 6 and Fig. 7 numerous 2-days simulations are performed. The reference control policy obtained through the GPDP algorithm using $\sigma=0.10$ is switched to other control strategies starting from the second day. It is again noticeable how the degradation in control performance, represented by BGLs falling outside the normoglycemic range, causes major surprise levels.

A special analysis for the fuzzy controller is presented in Fig. 8. In (a) the carbohydrate intake described in Table 1 is reduced in a 50%. This event is followed by lower BGL and respectively reduced insulin infusion rates to maintain the current control policy. But even when this represents a new balance inside the system, it does not represent a change in the behavior of the control policy and hence there exist no significant variations in the surprise levels. Both remaining curves present misconducts in the second half. In (b) the controller behavior is affected by retuning the membership functions for proportional, integral and derivative actions. In (c) insulin infusions are limited to 35 mU/L. This fact causes saturation in the manipulated variable and allows larger BGL, as represented by the two peaks in the surprise level.

4 Final Remarks

Based on *Bayesian surprise* and optimal action selection, a novel framework for control loop monitoring was presented. A unique property of the proposed approach is that can be used for automated supervision of Real-time Optimization (RTO) and dynamic optimization schemes working in uncertain environments.

Since any deviation from a reference behavior is correlated to a measurement of surprise, this framework may be applied to a variety of real time socio-technical systems. Some related fields on which Bayesian surprise could be useful include salient object detection for computer vision applications (Vig et al., 2012) and performance assessment of industrial controllers (Ghraizi et al., 2009). Also, Bayesian surprise may be successfully used for intensity, color, and motion features identification in image and video processing (Guo et al., 2010) as well as in sudden path deviation or lane change on autonomous vehicle systems (Chew et al. 2008).

References

1. Jaremko, J., Rorstad, O.: Advances towards the implantable artificial pancreas for treatment of diabetes. *Diabetes Care*. 21, 444–50 (1998)
2. Todorov, E.: Efficient computation of optimal actions. *PNAS*. 106, 28, 11478–11483 (2009)
3. Kárný, M., Guy, T. V.: Fully probabilistic control design. *Systems & Control Letters*. 55, 4, 259-265 (2006)
4. Plastino, A. R., Miller, H. G.: Minimum Kullback entropy approach to the Fokker-Planck equation. *Physical Review E*. 56, 4, 3927-3934 (1997)
5. Günel, S.: An information-theoretic approach to nonlinear systems: In the view of the Fokker-Planck-Kolmogorov formalism. LAP Lambert Academic Publishing, Germany (2010)
6. Majda, A., Kleeman, R., Cai, D.: A mathematical framework for quantifying predictability through relative entropy. *Methods and Applications of Analysis*. 9, 3, 425–444 (2002)
7. Ghraizi, R., Martínez, E., de Prada, C.: Control Loop Performance Monitoring using the Permutation Entropy of Error Residuals. *Proc. of ADCHEM*, ID199, Istanbul (2009)
8. Baldi, P., Itti, L.: Of bits and wows: A Bayesian theory of surprise with applications to attention. *Neural Networks*. 23, 649-666 (2010)
9. Owens, L., Doyle, J.: Performance Monitoring Of Diabetic Patient Systems. 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society. Istanbul, Turkey (2001)
10. Acikgoz, S. U., Diwekar, U. M.: Blood glucose regulation with stochastic optimal control for insulin-dependent diabetic patients. *Chemical Engineering Science*. 65, 3, 1227-1236 (2010)
11. Gelman, A., Carlin, J. B., Stern, H. S., Rubin, D. B.: Bayesian data analysis. Chapman and Hall, London (1995)
12. Deisenroth, M., Rasmussen, C., Peters, J.: Gaussian process dynamic programming. *Neurocomputing*. 72, 1508–1524 (2009)
13. De Paula, M., Martinez, E.: Probabilistic optimal control of blood glucose under uncertainty. Accepted as an oral presentation to the 22nd European Symposium on Computer Aided Process Engineering, London, June 17-21 (2012)
14. Farmer, T. G., Edgar, T. F., Peppas, N. A.: Effectiveness of intravenous infusion algorithms for glucose control in diabetic patients using different simulation models. *Ind. Eng. Chemical*. 48, 9, 4402–4414 (2009)
15. Chee, F., Fernando, T., Savkin, A., van Heeden, V.: Expert pid control system for blood glucose control in critically ill patients. *IEEE Trans. Inf. Technol. Biomed.* 7, 4, pp. 419–425 (2003)
16. Susanto-Lee, R., Fernando, T., Sreeram, V.: Simulation of Fuzzy-Modified Expert PID Algorithms for Blood Glucose Control. 10th Intl. Conf. on Control, Automation, Robotics and Vision, p. 1583 - 1589 (2008)
17. Vig, E., Dorr, M., Martinecz, T., Barth, E.: Intrinsic Dimensionality Predicts the Saliency of Natural Dynamic Scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34, 6, pp.1080-1091 (2012)
18. Guo, C., Zhang, L.: A Novel Multiresolution Spatiotemporal Saliency Detection Model and Its Applications in Image and Video Compression. *IEEE Transactions on Image Processing*, 19, 1, pp.185-198 (2010)
19. Chew, T., Adams, M., Ibañez-Guzmán, J.: A Relative Information Metric For Vehicle Following Systems. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nice, France, September 22-26 (2008)