

PEPPER: Patient Empowerment Through Predictive Personalised Decision Support

Pau Herrero,¹ Beatriz López² and Clare Martin³

Abstract. PEPPER is a newly-launched three-year research project, funded by the EU Horizon 2020 Framework. It will create a portable personalised decision support system to empower individuals on insulin therapy to self-manage their condition. PEPPER employs Case-Based Reasoning to advise about insulin bolus doses, drawing on various sources of physiological, lifestyle, environmental and social data. It also uses a Model-Based Reasoning approach to maximise users' safety. The system will be integrated with an unobtrusive insulin patch pump and has a patient-centric development approach in order to improve patient self-efficacy and adherence to treatment.

1 INTRODUCTION

Type 1 diabetes (T1D) is a chronic disease caused by an autoimmune destruction of the pancreatic beta cells. This leaves the body unable to produce the insulin needed to regulate blood glucose levels. The condition is usually controlled through multiple daily injections (MDI) of insulin to mimic the natural insulin secretion of a healthy pancreas. Alternatively, some people are treated with continuous subcutaneous insulin infusion (CSII) via a wearable pump. In both cases the size of each insulin doses are chosen by the individual.

Decision support tools exist to support this process, such as insulin bolus calculators that use simple mathematical formulae based on metabolic parameters (i.e. insulin-to-carbohydrate ratio and insulin sensitivity factor) and an estimation of the active insulin from previous doses. Such tools are integrated into most insulin pumps [9], and some glucose meters. There is also an increasing adoption of decision support tools implemented on mobile devices [10], often in conjunction with remote data storage in the cloud, though few are approved by regulatory bodies such as the FDA. Some gather inputs via wearable sensors (i.e. continuous glucose monitors), but most of them rely on manual input. In practice, the latter are rarely used because most people with T1D find the process tedious and refuse to interact with such systems [10, 2]. Hence a guiding design principle for PEPPER is that wherever possible data is collected automatically, via wearable technology. The information collected by the sensors is managed by a Case-Based Reasoning (CBR) module to provide personalised insulin recommendations, while a second Model-Based Reasoning (MBR) module is used to maximise users' safety.

2 SYSTEM OVERVIEW

The PEPPER system shown in Figure 1 offers insulin dosing advice that is highly adaptive to the insulin needs of individuals by using

a CBR approach. It also guarantees individuals' safety by means of a MBR approach that includes predictive glucose alarms, automatic insulin suspension, carbohydrate recommendations and fault diagnosis. PEPPER offers a dual architecture to cater for both MDI or CSII treatment, the latter via the unobtrusive Cellnovo patch-pump (Cellnovo Ltd., UK). In both cases, the patient periodically wears a continuous glucose monitor (CGM) used to automatically evaluate glucose outcomes. An activity monitor, such the one integrated in the Cellnovo pump or a commercially available one (e.g. Fitbit), is included to determine physical activity automatically. Data from a capillary blood glucose meter is periodically gathered to calibrate the CGM or to be used in case CGM data is not available. Additional data such as food intake, alcohol consumption, hormonal cycles are input through the user interface of the handheld unit (smartphone or Cellnovo handset). All inputs are then fed to the CBR engine on the handheld unit, and used to calculate the corresponding insulin dose. The dose is then displayed for the user to accept or decline. If the recommendation is accepted, the unit wirelessly sends the corresponding command to the insulin pump, or the user manually injects the bolus using an insulin pen. In addition, the safety module triggers alarms to alert the user about predicted hypo- and hyperglycaemic events. In the case of impending hypoglycaemia, the system also recommends a personalised amount of carbohydrates to consume to eliminate hypoglycaemia and avoid rebound hyperglycaemia. It also suspends insulin delivery for pump users when glucose levels are forecast to be too low. If potentially dangerous events are not properly addressed by the subject, automatic alarms can be sent via an SMS service to the expert team and selected carers. When network connectivity is available, the handheld unit sends the recorded data to a remote secure server. Data is presented in meaningful visualisations and analysed periodically to find non-optimal glucose patterns.

2.1 Case-Based Reasoning for Insulin Dosing

Case-Based Reasoning (CBR) is a consolidated artificial intelligence technique, extensively applied in medicine, that tries to solve newly encountered problems by applying solutions learned from similar problems encountered in the past. In CBR, past situations are stored in *cases*, which represent knowledge related to the various aspects of the situation. The CBR cycle consists of four steps: *Retrieve* the most similar case or cases; *Reuse* the information in that case to solve the problem; *Revise* the proposed solution; *Retain* the parts of this experience likely to be useful for future problem solving [1].

The first project to use CBR to recommend changes in insulin therapy for T1D management was the T-IDDM project [3], where it was integrated with rule-based reasoning and a probabilistic model of the effects of insulin on blood glucose levels. More recently, the IDSDM

¹ Imperial College London, email: pherrero@imperial.ac.uk

² University of Girona, email: beatriz.lopez@udg.edu

³ Oxford Brookes University, email: cemartin@brookes.ac.uk




Figure 1. PEPPER architecture.

project [11] used CBR as the primary reasoning modality in a decision support tool for patients on insulin pump therapy, and introduced other factors into the calculations, such as life events that can influence blood glucose levels. However, both projects were intended for use by clinicians as opposed to the individuals with diabetes.

In PEPPER, the CBR cycle is divided into two parts: the local and remote. The local part runs on the handheld unit and the remote part on a server. Both parts contain a case-base and periodically the local case-base is synchronised with the remote case-base. The evaluation step of the CBR cycle occurs on the server and requires approval by an expert clinician before a new case is incorporated to the case-base. The CBR parameters include CGM and capillary glucose data, physical activity, time, location, basal insulin, hormone cycle, stress, alcohol, meal composition, and sleep. Most of these parameters are automatically collected (or calculated) by the handset unit. Exceptions include alcohol consumption, meal composition and hormone cycles, which need to be manually inputted. A prototype version of the algorithm has already been implemented and successfully tested *in silico* [8] and in subsequent pilot studies [12]. PEPPER builds on this prototype and furthers improves it by including more parameters and automatising their recording.

2.2 Model-Based Reasoning for Safety

Model-Based Reasoning (MBR) is defined as the interaction of observation and prediction [5]. On the one hand, there is the actual system (e.g T1D subject) whose behaviour can be observed; on the other hand, there is the model of the system from which predictions (e.g. glucose levels) can be made. Assuming that the models are correct, any discrepancy found between observations and predictions are defaults on the device (e.g. CGM or pump fault). MBR techniques have been previously proposed in the context of diabetes technology to constrain insulin delivery by an artificial pancreas [4], predict hypoglycaemic events [6] and detect CGM and insulin pump faults [7]. PEPPER leverages these techniques to build a system that guarantees safety of the user at any time. In addition, it incorporates an adaptive carbohydrate recommender system to prevent hypoglycaemic events.

3 CONCLUSION

The PEPPER system provides a portable personalised decision support system for insulin dosing that combines data from multiple sources such as body-worn sensors and manual inputs. The Case-Based Reasoning module is designed to provide a personalised insulin dose which adapts over time. A Model-Based Reasoning module is designed to maximise safety through prediction of adverse events and the detection of faults. PEPPER is being developed using a patient-centric approach in order to improve patient self-efficacy and adherence to treatment. The software development will adhere to international standards including those that apply to security and interoperability. The final system will be tested *in silico* before being clinically validated over a 6-month non-randomised open-label ambulatory trial.

ACKNOWLEDGEMENTS

Thanks to the PEPPER team: A. Aldea, D. Brown, D. Duce, JM Fernández-Real, P. Georgiou, J. González López, R. Harrison, B. Innocenti, J. Masoud, L. Nita, N. Oliver, P. Pesl, R. Petite, M. Reddy, J. Shapley, F. Torrent, C. Toumazou, JP Vargheese, and M. Waite. This project has received funding from the EU Horizon 2020 research and innovation programme under grant agreement No 689810.

REFERENCES

- [1] Agnar Aamodt and Enric Plaza, 'Case-based reasoning: Foundational issues, methodological variations, and system approaches', *AI communications*, 7(1), 39–59, (1994).
- [2] Eirik Årsand, Dag Helge Frøisland, Stein Olav Skrøvseth, Taridzo Chomutare, Naoe Tatara, Gunnar Hartvigsen, and James T Tufano, 'Mobile health applications to assist patients with diabetes: lessons learned and design implications', *Journal of diabetes science and technology*, 6(5), 1197–1206, (2012).
- [3] Riccardo Bellazzi, Cristiana Larizza, Stefania Montani, Alberto Riva, Mario Stefanelli, Giuseppe d'Annunzio, R Lorini, Enrique J Gómez, E Hernando, Eulalia Brugués, et al., 'A telemedicine support for diabetes management: the t-iddm project', *Computer methods and programs in biomedicine*, 69(2), 147–161, (2002).
- [4] Ellingsen C., Dassau E., Zisser H., and et al., 'Safety constraints in an artificial pancreatic cell: An implementation of model predictive control with insulin on board', *Journal of diabetes science and technology*, 3(3), 536–544, (2009).
- [5] Randall Davis and Walter C. Hamscher, 'Model-based reasoning: Troubleshooting', in *Exploring Artificial Intelligence*, pp. 297–346. Morgan Kaufmann Publishers Inc., (1988).
- [6] RA Harvey, E Dassau, H Zisser, DE Seborg, L Jovanovic, and Doyle FJ., 'Design of the health monitoring system for the artificial pancreas: Low glucose prediction module.', *Journal of Diabetes Science and Technology*, 6(6), 1345–1354, (2012).
- [7] P Herrero, R Calm, J Veh, J Armengol, P Georgiou, N Oliver, and Tomazou C., 'Robust fault detection system for insulin pump therapy using continuous glucose monitoring.', *Journal of Diabetes Science and Technology*, 6(5), 1131–41, (2012).
- [8] Pau Herrero, Peter Pesl, Monika Reddy, Nick Oliver, Pantelis Georgiou, and Christofer Toumazou, 'Advanced insulin bolus advisor based on run-to-run control and case-based reasoning', *Biomedical and Health Informatics, IEEE Journal of*, 19(3), 1087–1096, (2015).
- [9] David C Klonoff, 'The current status of bolus calculator decision-support software', *J Diabetes Sci Technol*, 6(5), 990–994, (2012).
- [10] David C Klonoff, 'Telemedicine for diabetes current and future trends', *Journal of diabetes science and technology*, 10(1), 3–5, (2016).
- [11] Cindy Marling, Jay Shubrook, and Frank Schwartz, 'Case-based decision support for patients with type 1 diabetes on insulin pump therapy', in *Advances in Case-Based Reasoning*, 325–339, Springer, (2008).
- [12] Monika Reddy, Peter Pesl, Pau Herrero, and et al., 'Clinical safety and feasibility of the advanced bolus calculator for type 1 diabetes based on case-based reasoning: a 6-month randomised single-arm pilot study.', *Diabetes Technol Ther* 2016, (2016).