DIABETES TECHNOLOGY & THERAPEUTICS
Volume 18, Number 10, 2016

DOI: 10.1089/dia.2015.0266

DTT

ORIGINAL ARTICLE

Physiology-Invariant Meal Detection for Type 1 Diabetes

James Weimer, PhD,[1,*] Sanjian Chen, PhD,[1,*] Amy Peleckis, NP,[2]
Michael R. Rickels, MD, MS,[2] and Insup Lee, PhD[1]

Abstract

Background: Fully automated artificial pancreas systems require meal detectors to supplement blood glucose level regulation, where false meal detections can cause unnecessary insulin delivery with potentially fatal consequences, and missed detections may cause the patient to experience extreme hyperglycemia. Most existing meal detectors monitor various measures of glucose rate-of-change to detect meals where varying physiology and meal content complicate balancing detector sensitivity versus specificity.
Methods: We developed a novel meal detector based on a minimal glucose–insulin metabolism model and show that the detector is, by design, invariant to patient-specific physiological parameters in the minimal model. Our physiological parameter-invariant (PAIN) detector achieves a near-constant false alarm rate across all individuals and is evaluated against three other major existing meal detectors on a clinical type 1 diabetes data set.
Results: In the clinical evaluation, the PAIN-based detector achieves an 86.9% sensitivity for an average false alarm rate of two alarms per day. In addition, for all false alarm rates, the PAIN-based detector performance is significantly better than three other existing meal detectors. In addition, the evaluation results show that the PAIN-based detector uniquely (as compared with the other meal detectors) has low variance in detection and false alarm rates across all patients, without patient-specific personalization.
Conclusions: The PAIN-based meal detector has demonstrated better detection performance than existing meal detectors, and it has the unique strength of achieving a consistent performance across a population with varying physiology without any individual-level parameter tuning or training.

Background

Type 1 diabetes (T1D) affects ~1.25 million people in the United States and 5 million Americans are expected to have T1D by 2050.[1] T1D patients depend on daily insulin therapy to control glucose levels to avoid numerous long-term complications associated with hyperglycemia.[2] Meal carbohydrates cause significant disturbance to one's glucose level in T1D, making it critical to cautiously plan insulin injections around meal times to avoid postprandial hyperglycemia and subsequent postcorrection hypoglycemia. Artificial pancreas (AP) systems[3–5] aim to regulate the glucose level by automatically delivering insulin and free T1D patients from the cognitive burden of frequent glucose monitoring, carb counting, and insulin dosing decision making.

A key challenge of meal-time glycemic control is the sensing and action delays: the glucose level starts to rise a certain time after the onset of a meal and there is a delay between the injection of insulin and the action of insulin to dispose of glucose. To cope with this challenge, the AP systems need meal information either by announcement[6] or by meal detection.[7–13] This work concerns accurate and timely meal detection—that is, detecting whether carbohydrates have been ingested in the recent past. Accurate meal detection not only serves as the first step toward meal estimation (i.e., estimating the amount of carbohydrates ingested) but can also be employed by meal estimation algorithms as a safety backup, especially in situations where user input is erroneous.[10] The problems of carbohydrate estimation and insulin bolus calculation[13] are not considered herein and left as future work.

A meal detector aims to identify, in real-time, carbohydrate ingestion based on continuous glucose monitor (CGM) readings. Several meal-detection strategies already exist in the literature. Dassau et al. propose a voting-based meal detector that tracks the glucose rate-of-changes (RoCs) estimated by

[1]Department of Computer and Information Science, University of Pennsylvania, Philadelphia, Pennsylvania.
[2]Division of Endocrinology, Diabetes and Metabolism, Perelman School of Medicine, University of Pennsylvania, Philadelphia, Pennsylvania.
*These authors have contributed equally to this work and should be considered ''co-first authors.''

different methods including Kalman Filtering and announces a meal when three out of the four RoC estimates cross prespecified thresholds.[10] Using similar Kalman Filtering techniques, Lee and Bequette develop a meal detector that announces a meal based on RoCs crossing thresholds and estimates the meal size by feeding the filtered glucose RoCs into a finite response filter.[14] Harvey et al. recently proposed a meal-detection algorithm that announces meals based on a two-stage CGM filtering process and an RoC criteria.[11] Cameron et al. developed a meal-detection algorithm that uses a simple glucose model to match the RoC of the CGM readings to temporal trajectories assuming both no-meal and meal scenarios.[7]

All the aforementioned meal detectors require identifying patient-specific parameters (e.g., insulin sensitivity and insulin diffusion rate), most of which vary with time. Due to the inherent physiological dependency, the RoC-based detectors may suffer from high false positives, considering that nonmeal disturbance factors may also cause significant glucose fluctuations (e.g., exercising,[15] stress,[16] and depletion of insulin-on-board[17]). In addition, the trajectory-matching meal detector has a long average detection delay.[18] As an alternative, recent work by Turksoy et al. simultaneously aims to estimate physiological variables and model parameters to provide accurate meal detection and estimation; however, no guarantees are provided that the physiological parameter estimates converge to their true value.[13] Quick and reliable meal detection is critical for the AP systems: false detections can lead to unnecessary insulin delivery that may trigger life-threatening hypoglycemia; missed detections or significant detection delays can leave the patient with marked postprandial hyperglycemia.

This article presents a novel meal-detection algorithm that is based on a commonly accepted minimal glucose physiological model and is ''invariant'' to individual physiological parameters—that is, it achieves a near-constant false alarm rate (CFAR) across the population without needing individual tuning. We compare our meal detector with three published meal-detection techniques.[10,11,14] Evaluations on a real T1D clinical data set collected through the Rodebaugh Diabetes Center of the University of Pennsylvania Health System (UPHS) show that our detector outperforms (often significantly) other detectors in terms of detection rate and false-positive rate, while also having significantly less performance variation between patients.

Methods

In this section, we present our meal detector based on a parameter-invariant (PAIN) design approach[25] used to achieve a CFAR. In many medical monitoring applications, including meal detection, unknown or uncertain physiology presents a fundamental challenge in generating mathematical models useful for detector design. The PAIN approach uses physiological models and trends to capture the effects of the unknown nuisance parameters, then establishes invariance to the nuisance parameters by projecting the measurements onto a space that is unaffected by the unknown parameters, mathematically known as a null space projection. The benefit of the PAIN approach is that the projected measurements will be the same, regardless of the patient's unknown physiology, allowing the design of powerful detectors that leverage this population-level consistency. The PAIN design approach has been successfully applied to various engineering applications with unknown parameters[19–21] and has recently been extended to medical monitor design.[22–25]

To present the meal detector, this section uses Figure 1 as a visual aid and references the Appendix for useful mathematics related to its formulation. In this work, we perform meal detection in two steps. First, at time k shown in Figure 1, we use a time window of measurements (denoted by w in Fig. 1) and comparatively test, using PAIN techniques,[25] two consecutive subwindows of time ending δ time steps before k (represented by d_0 and d_1 in Fig. 1) for the presence of a meal. The time steps correspond to a CGM sampling period of 1 min. Second, we sequentially filter the test decisions generated at each time step k to generate a threshold-based test for meal detection. The remainder of this section describes, in detail, the key components of the meal detector, namely, modeling glucose–insulin physiological trends, designing physiology-invariant tests, and filtering test decisions.

Modeling glucose–insulin physiological trends

Many models exist for describing glucose–insulin physiology, ranging from high-fidelity maximal models[26,27] to an assortment of low-fidelity minimal models.[28] For the PAIN technique to be useful, the chosen model must capture the general physiological trends that discriminate meal occurrence or absence. Thus, to capture the real-life scenario where the glucose level is measured at a subcutaneous site and carbohydrates enter plasma through a digestion pathway, we use a modified 5th-order linear Bergman model[28] augmented with minimal compartmental models that describe the subcutaneous insulin pathway[29,30] and meal carbohydrate digestion pathway.[31]

The complete augmented physiological model is a five-state linear system (provided in the Appendix, Equation 1) and contains several specific physiological parameters, for example, the insulin sensitivity,[28] the insulin diffusion rate,[30] and the time of maximum glucose appearance.[31] Identifying these parameters for each individual requires time-consuming tests in strictly controlled clinical settings, which may be inaccurate outside the controlled setting. Thus, a core element of our meal detector is the design of tests invariant to the unknown time-varying physiological parameters.

Designing physiology-invariant tests

Applying standard time-series analysis techniques,[32] we can write the CGM measurement model at time step k (as shown in Fig. 1), assuming meal window d_i for $i \in \{0, 1\}$, as $\boldsymbol{y_k} = \boldsymbol{H_{k,i}}\boldsymbol{\theta} + \boldsymbol{\sigma n}$, where $\boldsymbol{y_k}$ is a vector of the w CGM measurements, and $\boldsymbol{H_{k,i}}$ is a known matrix (defined in the Appendix, Equation 2) that relates how the CGM measurements are affected by the lumped-physiological parameters, $\boldsymbol{\theta}$. The value of $\boldsymbol{\theta}$ is a function of the unknown physiological variables.[a] In addition, $\boldsymbol{\sigma}$ represents an unknown uncertainty associated with a zero-mean noise, $\boldsymbol{n}$. Using the CGM measurement model, we can generate two invariant statistics, $t_0(\boldsymbol{y_k})$ and $t_1(\boldsymbol{y_k})$, as defined in Equation 3 in the Appendix.

[a]In this work, we omit the specific mapping of physiological variables to the lumped parameters as it is irrelevant in the design of PAIN detectors,[32] that is, designing tests invariant to the lumped parameters is equivalent to designing invariant to the underlying physiological variables.

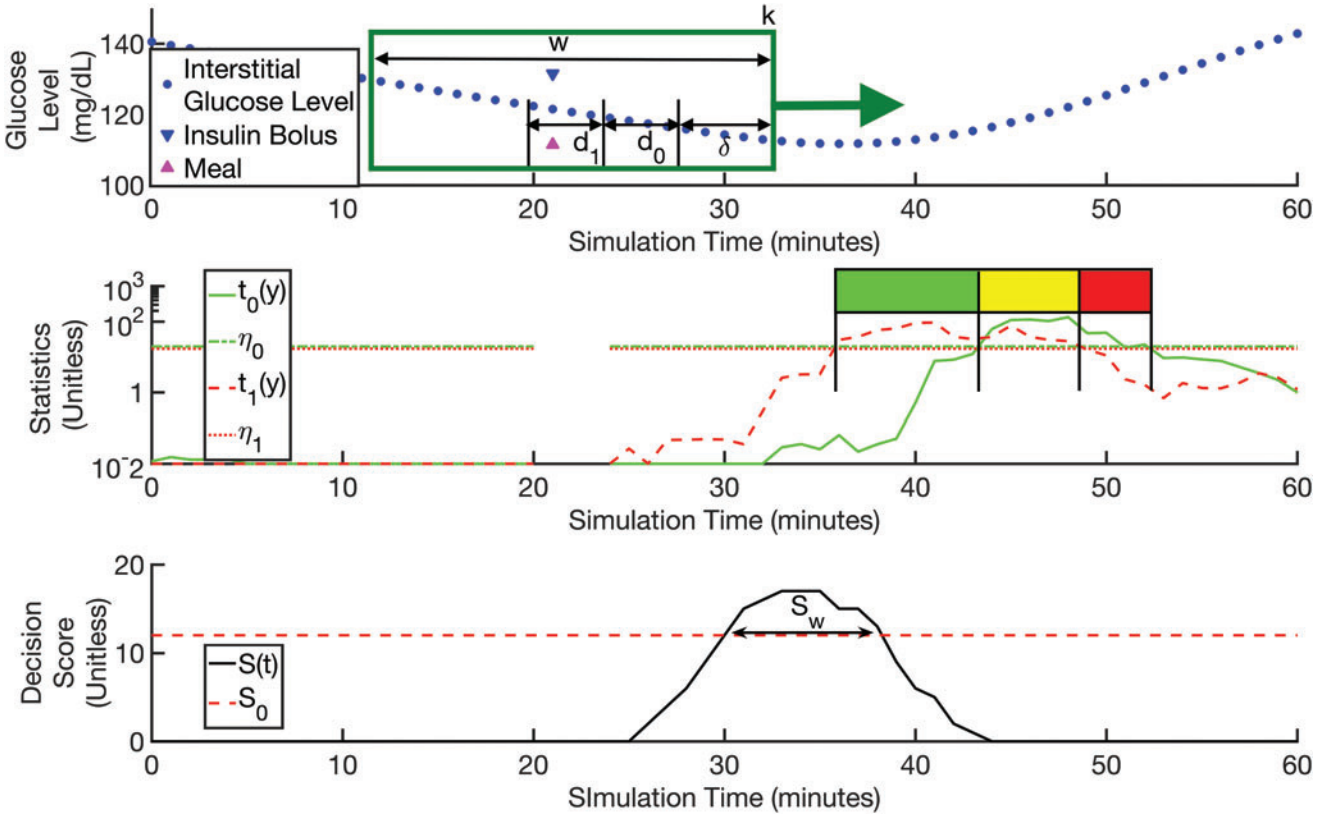

FIG. 1. A meal detection example of the PAIN-based detector. PAIN, parameter invariant.

In other words, for $t_0(\boldsymbol{y_k})$, we project the CGM measurements onto the null space of $\boldsymbol{H_{k,\theta}}$,[33] then generate $t_0(\boldsymbol{y_k})$ using the ratio of the remaining measurement energy in the space of $\boldsymbol{H_{k,1}}$ to the measurement energy not remaining in $\boldsymbol{H_{k,1}}$. The form of $t_0(\boldsymbol{y_k})$ is commonly referred to as an F-ratio in the signal processing and statistics literature,[32] and has the useful feature that its value is invariant to the noise level $\boldsymbol{\sigma}$ as well as the lumped-physiological parameters $\boldsymbol{\theta}$. In the context of this work, $t_0(\boldsymbol{y_k})$ represents the ratio of measurement energy aligned with (and only with) the meal effects of d_1 to the measurement energy that cannot be explained exclusively by meals within d_1. Comparing $t_0(\boldsymbol{y_k})$ to a threshold η_0, selected to achieve a specified probability of false alarm, generates a decision. Similarly, $t_1(\boldsymbol{y_k})$ is generated by first projecting the measurements onto the null space of $\boldsymbol{H_{k,1}}$, then generating an F-statistic using $\boldsymbol{H_{k,\theta}}$ and comparing with a threshold η_1.

The selection of the PAIN-based detector parameters, d_0, d_1, δ, and w, can significantly affect its performance. In this work, we select d_0 and d_1 to be 10 time steps, δ to be 20 time steps, and w to be 120 time steps. These values are chosen because they provided the best detection rates among the range of values evaluated. A discussion of the PAIN-based detector parameter effects is provided in the Appendix along with a detailed presentation of the test statistics.

Figure 1 illustrates how the PAIN-based detector[b] works on simulated scenarios generated by the FDA-accepted simulator.[27] In this figure, we use simulated data (as opposed to the real population data used for evaluation) to clearly illustrate our approach to meal detection; thus, the CGM measurements correspond to a 1-min sampling rate of the interstitial glucose level shown in Figure 1. The true meal happens around 22 min (the pink upper triangle in the top plot of Fig. 1). The PAIN-based detector works in a sliding-window manner: at time k, the detector run tests on the d_0 and d_1 windows using the past w CGM measurements; the relevant time windows at time k are scoped by the green box in Figure 1; the detector generates a decision at each time and the time windows (as highlighted in the green box) shift forward in time with the detector to generate sequential statistics (whose values are shown in the second and third subfigures in Fig. 1).

In Figure 1, as the d_0 window approaches the true meal event (the detector never knows when a meal actually happens and tests every time step), the statistic $t_1(\boldsymbol{y_k})$ (the red dashed line in the figure's middle graph) starts rising and becomes separated from $t_0(\boldsymbol{y_k})$ (the solid green line), indicating a meal is more likely to have occurred in d_0 than in d_1. Then as the detector moves further ahead, the true meal enters the d_1 window, and $t_0(\boldsymbol{y_k})$ increases while $t_1(\boldsymbol{y_k})$ decreases, indicating that a meal is more likely in d_1 than in d_0. This sequential rise and fall of the statistics $t_0(\boldsymbol{y_k})$ and $t_1(\boldsymbol{y_k})$ are leveraged to design a sequential test.

Filtering sequential test decisions

To leverage the structured sequential rise and fall of the statistics, we design an algorithm that generates a cumulative decision score based on the statistics $t_0(\boldsymbol{y_k})$ and $t_0(\boldsymbol{y_k})$ and corresponding thresholds η_0 and η_1. The statistics have the useful property that an increasingly positive $t_0(\boldsymbol{y_k})$ implies an

[b]For aesthetics, the values of d_0, d_1, δ, and w shown in Figure 1 are not the same as the values used in the evaluation.

TABLE 1. SCORE ACCUMULATION RULES FOR $S(j)$ AT k

	$t_0(\boldsymbol{y_k}) > \eta_0$	$t_0(\boldsymbol{y_k}) \leq \eta_0$
$t_1(\boldsymbol{y_k}) > \eta_1$	Meal in d_0 or d_1 $+ t_1(\boldsymbol{y_k})$ to $j \in d_0$ $+ t_0(\boldsymbol{y_k})$ to $j \in d_1$	Meal in d_0 $+ 2 * t_1(\boldsymbol{y_k})$ to $j \in d_0$
$t_1(\boldsymbol{y_k}) \leq \eta_1$	Meal in d_1 $+ 2 * t_0(\boldsymbol{y_k})$ to $j \in d_1$	No Meal Do not change $S(j)$

increasing likelihood that a meal has occurred in the window d_1 (and vice versa). Thus, the algorithm generates an S-score, $S(j)$, for each time step j (assuming $S(j)$ is initialized to zero) and accumulates S-scores according to the rules in Table 1, where a larger S-score indicates a higher confidence in the occurrence of a meal.

At every step, when the detector claims a meal occurs in window d_0, we add $2 * t_1(\boldsymbol{y_k})$ to $S(j)$ for each time step j in the d_0 window; similarly, if the detector claims a meal occurs in d_1, we add $2 * t_0(\boldsymbol{y_k})$ to $S(j)$ for each time step in the d_1 window. If it is likely that a meal was in both windows, then we add $t_1(\boldsymbol{y_k})$ to $S(j)$ for each time step in d_0 and similarly we add $t_0(\boldsymbol{y_k})$ to $S(j)$ of each time step in d_1. Note that we drop the factor of 2 in the increments when both residual statistics, $t_0(\boldsymbol{y_k})$ and $t_1(\boldsymbol{y_k})$, are positive, thus weakening the confidence of a meal happening in any individual window. If both residual statistics are negative, then neither d_0 nor d_1 is likely to contain a meal; thus, no score accumulation occurs.

The score accumulation rules are highlighted in Figure 1: each colored region corresponds to the rule in Table 1 that applies in that region. In a typical positive meal-detection scenario, one should first see the green region (corresponding to the d_0 window) approaches the meal event, followed by the yellow region as the meal event transitions from d_0 to d_1, and finally see the red region, after which a peak in the $S(t)$ curve emerges, indicating that the detector makes a series of decisions at sequential time steps that all point to the same meal time region where the $S(t)$ peak emerges.

The magnitude of $S(t)$ corresponds to our confidence in a meal occurring at time t. To trigger an alarm (indicating a meal has occurred), we use two design parameters, a threshold S_0 and a minimum width S_w; a peak is characterized by at least S_w consecutive $S(j)$'s that are above S_0. At each time step, the detector raises a meal alarm if a new $S(t)$ peak emerges. Parameters S_0 and S_w can be tuned to achieve different detection performance: smaller S_0 and S_w result in higher

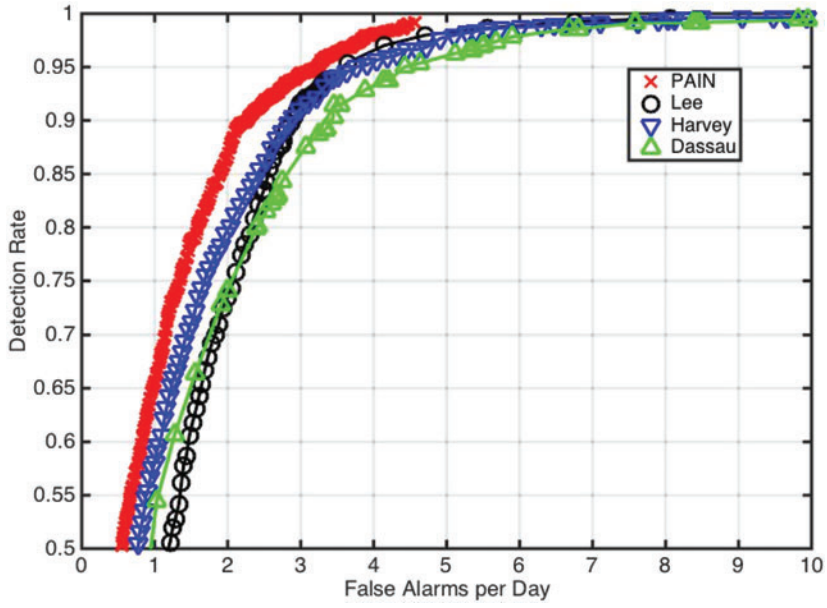

FIG. 2. The receiver operating characteristic curves of the four meal detectors on the clinical data set used for analysis.

Table 2. Operating Points of the Four Detectors on the Clinical Data set Used for Analysis

Detector	*Detection rate (%)*	*False alarms per day*
PAIN	86.9	2.01
Dassau[10]	74.1	1.99
Lee[14]	73.4	1.99
Harvey[11]	79.4	1.97

PAIN, parameter invariant.

sensitivity but more false alarms. We note that there are a few steps delay between the actual meal time and the $S(t)$ peak, as shown in Figure 1. This delay phenomenon is related to the physiological fact that there is a delay from the onset of eating to when the CGM reading starts changing: in the maximal model,[26,27] meal carbohydrates have to pass several digestion compartments before affecting the plasma glucose.

Experimental Results

This section presents the evaluation results of the PAIN-based detector against three existing meal-detection algorithms: the Dassau (et al.) detector,[10] Harvey (et al.) detector,[11] and Lee (and Bequette) detector.[14] We evaluate the detectors using a clinical data set collected by the Rodebaugh Diabetes Center of the UPHS. The clinical data set includes 5-min CGM readings from 61 T1D patients (mean ± standard deviation of age: 45.7 ± 15.3 years; mean ± standard deviation of body weight: 79.2 ± 21.9 kg; average duration of monitoring 17 days).

Each patient counts the carbohydrates in the meal and then inputs that information into his or her insulin pump. The insulin pump will then provide a suggested meal bolus that the patient can accept or override. Since patient reporting of meal time (i.e., the time when the patient inputs the information into the pump) is error prone, we consider any meal alarm within 2 h of the reported meal event to be a correct detection (and a false alarm otherwise). In the event that a meal detector alarms within 30 min of a correction bolus (i.e., nonmeal bolus), we omit this alarm from our false alarm (specificity) analysis since the alarm corresponds to a situation (e.g., unreported meal, ''meal-like'' critical event, worsening of long-lasting hyperglycemia) requiring an insulin bolus.

We ran the PAIN-based detector, the Dassau detector, Harvey detector, and Lee detector on the same glucose measurements from the 61 patients. Each of the four meal detectors has a set of configurable parameters, for example, the threshold S_0 of the PAIN-based detector and RoC thresholds of the RoC-based detectors. We systematically explore the combinations of each detector's parameters and get their best detection performance.

A receiver operating characteristic (ROC) curve represents the detection rate and false alarm rate of a detector under different configurations. Figure 2 shows the ROC curves of the four detectors. Table 2 lists the operating point of each detector closest to two false alarms per day. This operating point was chosen to compare the relative detection performance (sensitivity) of each detector for a chosen specificity. For all false alarm rates (specificities), the PAIN-based detector has superior performance, where for the operating point reported in Table 2 has an 86.9% sensitivity (meaning that the detector correctly detects 86.9% of meal events within 2 h of the reported meal time) and 2.01 false alarms per day. We note that the evaluation results of the Harvey detector are comparable with those reported in the original publication.[11,c]

Figure 3 compares the performance variability, in terms of false alarm and detection rates, of each meal detector on different patients in the data set. These results provide a measure of the consistency of detection performance at the individual level, that is, whether a detector can perform particularly bad on any patient. The duration of glucose monitoring varies across patients in the data set. Eight patients are excluded from the individual-level analysis because their data contains fewer than 10 reported meals. Over the remaining 53 patients, the PAIN-based detector detects at least 55% of all reported meals and never has a false alarm rate greater than 3.7 false alarms per day. In sharp contrast to the PAIN-based detector, all other three detectors miss significantly more meals (both on average and worst case), and have false alarm rates with higher variances and higher worst-case values.

Discussion

The experimental evaluation shows that the PAIN-based detector significantly improves the detection performance when compared with the other three detectors. For example, compared with the Harvey detector, the PAIN-based detector reduces the number of missed detections by 36% without increasing the false alarm rate in the experimental evaluation. The performance distribution over the patients validates the unique strength of the PAIN-based detector: it is designed to be ''invariant'' to differences in patients' physiological parameters and thereby achieves low variance detection performance across a real population. This unique feature of the PAIN-based detector is critical to the AP: A meal detector that frequently misses true meal events on some patients could result in postprandial hyperglycemia and possibly subsequent post-correction hypoglycemia.

In theory, the performance of the RoC-based detectors may be further improved by carefully tuning the detector parameters for each individual. However, such tuning process may require frequent clinic visits because patients' physiological characteristics change over time. Moreover, even with parameter tuning, the RoC-based meal detectors have their fundamental limitation because the postmeal glucose rise rate depends on many other factors such as the nutrition composition of meals[34] and insulin-on-board,[35] which cannot be mitigated by simply tuning the threshold parameters. In contrast, the experimental evaluation demonstrates that the PAIN-based detector is able to achieve low-variance performance without any individual-level parameter tuning.

Conclusion

This article proposes a PAIN-based meal detector that is based on a physiological model. The clinical data set evaluation demonstrates that the PAIN-based meal detector has better detection performance than three existing meal detectors. In addition, the evaluation results validate that

[c]Clinical evaluation results of sensitivity versus specificity of the Dassau detector and the Lee detector are not available in their original publications[10,14]: the Dassau detector is evaluated in a single-meal test[10] and the Lee detector is evaluated using a simulator.[14]

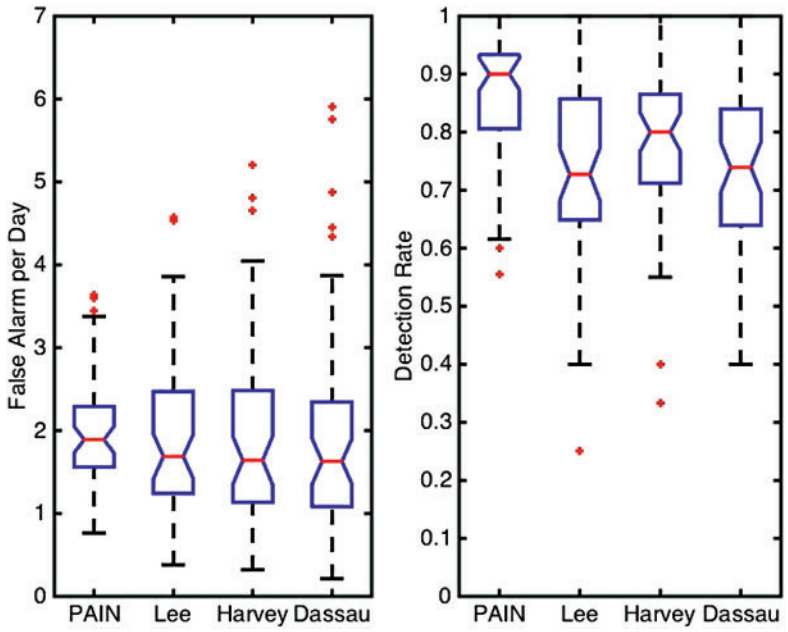

FIG. 3. Distributions of the four detectors' false alarm and detection rates on different patients in the clinical data set used for analysis. The '+' marks corresponds to outliers in patient performance for each detector.

the PAIN-based detector has the unique strength of achieving highly consistent performance across a real population, with varying physiology, without any individual-level parameter tuning. Although accurate and timely meal detection remains an open problem, the performance improvement provided by the PAIN approach suggests it is better suited, compared with other approaches, to provide meal alerts to a closed-loop controller and/or to act as a safety backup for meal information provided by a user or other meal detectors.

Acknowledgments

This research was supported, in part, by NSF grant CNS-1035715 (to I.L.), in part, by NIH grant R01- DK091331 (to M.R.R.), and, in part, by the DGIST Research and Development Program of the Ministry of Science, ICT, and Future Planning of Korea (CPS Global Center). The authors would like to acknowledge support from the Human Metabolism Resource of the Institute for Diabetes, Obesity & Metabolism at the University of Pennsylvania.

Author Disclosure Statement

No competing financial interests exist.

References

1. Centers for Disease Control and Prevention, et al.: National diabetes statistics report: estimates of diabetes and its burden in the united states, 2014. Atlanta, GA: US Department of Health and Human Services, 2014.
2. Klein R: Hyperglycemia and microvascular and macrovascular disease in diabetes. Diabetes Care 1995;18:258–268.
3. Bruttomesso D, Farret A, Costa S, et al.: Closed-loop artificial pancreas using subcutaneous glucose sensing and insulin delivery and a model predictive control algorithm: preliminary studies in Padova and Montpellier. Diabetes Sci Technol 2009;3:1014–1021.
4. Cobelli C, Renard E, Kovatchev B: Artificial pancreas: past, present, future. Diabetes 2011;60:2672–2682.
5. Meece J: The artificial pancreas where we are, where we're going. AADE Pract 2015;3:42–44.
6. Magni L, Raimondo DM, Bossi L, et al.: Model predictive control of type 1 diabetes: an in silico trial. Diabetes Sci Technol 2007;1:804–812.
7. Cameron F, Niemeyer G, Buckingham BA: Probabilistic evolving meal detection and estimation of meal total glucose appearance. J Diabetes Sci Technol 2009;3:1022–1030.
8. Chen S, Weimer J, Rickels M, et al.: Towards a model-based meal detector for type I diabetics. Medical Cyber-Physical Systems Workshop 2015, 2015. http://repository.upenn.edu/cis_papers/782 (accessed July 10, 2016).
9. Cobelli C, Man CD, Sparacino G, et al.: Diabetes: models, signals, and control. IEEE Rev Biomed Eng 2009;2:54–96.
10. Dassau E, Bequette BW, Buckingham BA, Doyle FJ: Detection of a meal using continuous glucose monitoring implications for an artificial β-cell. Diabetes Care 2008;31:295–300.
11. Harvey RA, Dassau E, Zisser H, et al.: Design of the glucose rate increase detector: a meal detection module for the health monitoring system. J Diabetes Sci Technol 20148: 307–320.
12. Lee H, Buckingham BA, Wilson DM, Bequette BW: A closed-loop artificial pancreas using model predictive control and a sliding meal size estimator. J Diabetes Sci Technol 2009;3:1082–1090.
13. Turksoy K, Samadi S, Feng J, et al.: Meal-detection in patients with type 1 diabetes: a new module for the multivariable adaptive artificial pancreas control system. IEEE J Biomed Health Inform 2016;20:47–54.
14. Lee H, Bequette BW: A closed-loop artificial pancreas based on model predictive control: human-friendly identification and automatic meal disturbance rejection. Biomed Signal Process Control 2009;4:347–354.
15. Grimm J: Exercise in type 1 diabetes. Nagi, Dinesh, ed. In: Exercise and Sport in Diabetes. New York: John Wiley & Sons; 2005:25–43.
16. Capes SE, Hunt D, Malmberg K, et al.: Stress hyperglycemia and prognosis of stroke in nondiabetic and diabetic patients a systematic overview. Stroke 2001;32:2426–2432.

17. Reichard P, Nilsson B-Y, Rosenqvist U: The effect of long-term intensified insulin treatment on the development of microvascular complications of diabetes mellitus. N Engl J Med 1993;329:304–309.
18. Guelfi KJ, Jones TW, Fournier PA: The decline in blood glucose levels is less with intermittent high-intensity compared with moderate exercise in individuals with type 1 diabetes. Diabetes Care 2005;28:1289–1294.
19. Weimer J, Ahmadi SA, Araujo J, et al.: Active actuator fault detection and diagnostics in hvac systems. Proceedings of the Fourth Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings. Toronto, Canada; 2012:107–114.
20. Weimer J, Varagnolo D, Johansson KH: Distributed model-invariant detection of unknown inputs in networked systems. International Conference on High Confidence Networked Systems. Philadelphia, PA; 2013:127–134.
21. Weimer J, Varagnolo D, Stankovic MS, Johansson KH: Parameter-invariant detection of unknown inputs in networked systems. Conference on Decision and Control. Florence, Italy; 2013:4379–4384.
22. Ivanov R, Weimer J, Simpao A, et al.: Early detection of critical pulmonary shunts in infants. Proceedings of the 6th ACM/IEEE International Conference on Cyber-Physical Systems, ICCPS'15. Seattle, Washington; 2015.
23. Roederer A, Weimer J, Dimartino J, et al.: Towards non-invasive monitoring of hypovolemia in intensive care patients. Medical Cyber-Physical Systems Workshop 2015, Seattle, Washington; 2015.
24. Roederer A, Weimer J, Dimartino J, et al.: Robust monitoring of hypovolemia in intensive care patients using photoplethysmogram signals. IEEE Engineering in Medicine and Biology Conference, EMBC'15, Milan, Italy; 2015.
25. Weimer J, Ivanov R, Roederer A, et al.: Parameter invariant design of medical alarms. IEEE Design Test 2015;32:9–16.
26. Dalla Man C, Micheletto F, Lv D, et al.: The UVA/PADOVA type 1 diabetes simulator: new features. J Diabetes Sci Technol 2014;8:26–34.
27. Kovatchev BP, Breton M, Dalla Man C, Cobelli C: In silico preclinical trials: a proof of concept in closed-loop control of type 1 diabetes. J Diabetes Sci Technol 2009;3:44–55.
28. Bergman RN, Ider YZ, Bowden CR, Cobelli C: Quantitative estimation of insulin sensitivity. Am J Physiol Endocrinol Metab 1979;236:E667.
29. Kobayashi T, Sawano S, Itoh T, et al.: The pharmacokinetics of insulin after continuous subcutaneous infusion or bolus subcutaneous injection in diabetic patients. Diabetes 1983;32:331–336.
30. Nucci G, Cobelli C: Models of subcutaneous insulin kinetics. a critical review. Comput Methods Programs Biomed 2000;62:249–257.
31. Gillis R, Palerm CC, Zisser H, et al.: Glucose estimation and prediction through meal responses using ambulatory subject data for advisory mode model predictive control. J Diabetes Sci Technol 2007;1:825–833.
32. Scharf LL: Statistical Signal Processing, volume 98. Reading, MA: Addison-Wesley, 1991.
33. Horn RA, Johnson CR, eds. Matrix Analysis. New York: Cambridge University Press, 1986.
34. Wolever TM, Bolognesi C: Prediction of glucose and insulin responses of normal subjects after consuming mixed meals varying in energy, protein, fat, carbohydrate and glycemic index. J Nutr 1996;126:2807–2812.
35. Ellingsen C, Dassau E, Zisser H, et al.: Safety constraints in an artificial pancreatic β cell: an implementation of model predictive control with insulin on board. J Diabetes Sci Technol 2009;3:536–544.

Address correspondence to:
Sanjian Chen, PhD
Department of Computer and Information Science
University of Pennsylvania
3330 Walnut Street, Levine 302
Philadelphia, PA 19104

E-mail: sanjian@cis.upenn.edu

Michael R. Rickels, MD, MS
Division of Endocrinology, Diabetes and Metabolism
Perelman School of Medicine
University of Pennsylvania
12-134 Smilow Center for Translational Research
3400 Civic Center Boulevard
Philadelphia, PA 19104

E-mail: rickels@mail.med.upenn.edu

(*Appendix follows* →)

Appendix

This appendix provides the mathematical details and supporting discussion for implementation of the parameter-invariant detector. Deriving the detector test statistics requires null space transformations, where the null space of an arbitrary matrix $\boldsymbol{X}$ is[33]

$$\langle \boldsymbol{X} \rangle^{\perp} = \{ v \,|\boldsymbol{X}v = 0 \}$$

and has an orthonormal basis transposed, $\boldsymbol{X}^{\perp}$, satisfying[33]

$$\boldsymbol{X}^{\perp} \in \{ \boldsymbol{V} \,|\forall v \in \langle \boldsymbol{X} \rangle^{\perp}, \exists x, \boldsymbol{V^T} x = v \wedge \boldsymbol{V V^T} = \boldsymbol{I} \},$$

where $\boldsymbol{V^T}$ denotes the transpose of matrix $\boldsymbol{V}$.[33] The following employs the mentioned notation to present, mathematically, the meal-detector test statistics implemented in this work.

For completeness, we begin by stating the augmented 5th-order linear Bergman model employed in this work,

$$\begin{bmatrix} \dot{G}(t) \\ \dot{m}(t) \\ \dot{g}(t) \\ \dot{I}(t) \\ \dot{I}_s(t) \end{bmatrix} = \begin{bmatrix} p_1 & 1 & 0 & p_2 & 0 \\ 0 & -t_G^{-1} & t_G^{-1} & 0 & 0 \\ 0 & 0 & -t_G^{-1} & 0 & 0 \\ 0 & 0 & 0 & -k_e & \frac{k_a}{t_G} \\ 0 & 0 & 0 & 0 & -k_a \end{bmatrix} \times \begin{bmatrix} G(t) \\ m(t) \\ g(t) \\ I(t) \\ I_s(t) \end{bmatrix} + \begin{bmatrix} p_2 \\ 0 \\ D_G(t) \\ 0 \\ u(t) \end{bmatrix}, \quad [1]$$

where $G, m, g, I,$ *and* I_s denote the physiological state for plasma glucose, plasma glucose appearance rate, digestive compartment glucose, plasma insulin, and subcutaneous insulin, respectively. The insulin bolus and meals are represented by u and D_G, respectively. All other variables represent unknown physiological parameters. For a complete discussion of the model in Equation 1, see our previous publication [8], section 4.2. Applying standard time-series discretization techniques, the model in Equation 1 can be written as a 5th-order discrete time system, assuming piecewise constant insulin and meal inputs.

For any sampling rate, we denote the k-th sample of the CGM measurement as x_k (as sampled from G) and the injected insulin bolus as u_k. We then write

$$\boldsymbol{y_k} = [x_k \ldots x_{k-w}]^T$$

$$\boldsymbol{F_k} = \begin{bmatrix} x_{k-1} & \cdots & x_{k-5} & u_{k-1} & \cdots & u_{k-4} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ x_{k-w-1} & \cdots & x_{k-w-5} & u_{k-w-1} & \cdots & u_{k-w-4} \end{bmatrix},$$

where we call $\boldsymbol{y_k}$ the *measurements* (representing a point in the *measurement space*) and we say $\boldsymbol{F_k}$ *spans* the measurement space affected by the insulin bolus and physiological dynamics. More importantly, each column of $\boldsymbol{F_k}$ corresponds to the effect, on the measurements, of an unknown lumped-physiological parameter. The mapping of the physiological parameters in Equation 1 to the lumped-physiological parameters is unimportant in implementing the PAIN-based meal detector and consequently omitted from this appendix.

Although $\boldsymbol{F_k}$ spans the measurement space affected by insulin bolus and physiological parameters, it does not (necessarily) span the effect of meals on the measurements. We capture the effect of meals within the hypothesized meal windows, d_0 and d_1 in Figure 1, respectively, as

$$\boldsymbol{G_0} = \begin{bmatrix} \boldsymbol{0}_{(\delta-4)} \times (d_0+4) \\ \boldsymbol{I}_{(d_0+4)} \\ \boldsymbol{0}_{(w-\delta-d_0)} \times (d_0+4) \end{bmatrix} \text{ and}$$

$$\boldsymbol{G_1} = \begin{bmatrix} \boldsymbol{0}_{(\delta+d_0-4)\times(d_1+4)} \\ \boldsymbol{I}_{(d_1+4)} \\ \boldsymbol{0}_{(w-\delta-d_0-d_1)} \times (d_1+4) \end{bmatrix},$$

where $\boldsymbol{0}_{(\mathrm{n})\times(\mathrm{m})}$ denotes an n-by-m matrix of all zeros and $\mathbf{I}_{(\mathrm{m})}$ corresponds to the m-dimensional identity matrix. We note that $\boldsymbol{G_i}$ has $d_i + 4$ columns (as opposed to d_i columns) since the effect of the most recent hypothesized meal (of unknown magnitude) within the d_i window affects measurements up to four time steps later. Thus, we say that $\boldsymbol{G_i}$ *spans* the measurement subspace affected by meals within d_i (according to the Bergman model).

We write

$$\boldsymbol{H_{k,i}} = [\boldsymbol{F_k} \ \boldsymbol{G_i}], i \in \{0, 1\} \quad [2]$$

and say that $\boldsymbol{H_{k,i}}$ spans the measurement subspace affected by the combined effect of parameters corresponding to the physiological dynamics, insulin bolus, *and* the meals within the d_i time window. Assuming a meal occurs exclusively within the time window d_i, then $\boldsymbol{y_k} = \boldsymbol{H_{k,i}\theta} + \boldsymbol{\sigma n}$ as described in the text.

To present the test statistics, we introduce intermediate variables

$$\boldsymbol{r}_{k,0} = \boldsymbol{H}^{\perp}_{\boldsymbol{k,0}} \boldsymbol{y}, \quad \boldsymbol{U_{k,0}} = \boldsymbol{H}^{\perp}_{k,0} \boldsymbol{G_1}$$

$$\boldsymbol{r}_{k,1} = \boldsymbol{H}^{\perp}_{\boldsymbol{k,1}} \boldsymbol{y}, \quad \boldsymbol{U_{k,1}} = \boldsymbol{H}^{\perp}_{k,1} \boldsymbol{G_0},$$

where $\boldsymbol{r}_{k,0}$ and $\boldsymbol{U_{k,0}}$ denote the projection of the measurements and projected meal effects for d_1 onto the nullspace of $\boldsymbol{H_{k,0}}$, respectively (and vice versa for $\boldsymbol{r}_{k,1}$ and $\boldsymbol{U_{k,1}}$). In other words, $\boldsymbol{r}_{k,0}$ and $\boldsymbol{U_{k,0}}$ denote the measurements and the effects of meals within d_1 that cannot be explained by physiological parameters, insulin bolus, and meals occurring within d_0. Consequently, to quantify whether the projected measurements and projected meal effects are significantly aligned,[30] we write test statistics, $t_i(\boldsymbol{y_k})$ for $i \in \{0, 1\}$, as

$$t_i(\boldsymbol{y_k}) = \frac{r_{k,i}^T\left(\boldsymbol{I} - \left(\boldsymbol{U}_{k,i}^{\perp}\right)^T\left(\boldsymbol{U}_{k,i}^{\perp}\right)\right)r_{k,i}}{r_{k,i}^T\left(\boldsymbol{U}_{k,i}^{\perp}\right)^T\left(\boldsymbol{U}_{k,i}^{\perp}\right)r_{k,i}}. \quad [3]$$

For $t_0(\boldsymbol{y_k})$, the numerator denotes the magnitude of the projected measurements aligned with (i.e., in the subspace of) the projected meal effects of d_1, whereas the denominator represents the energy of the projected measurements that cannot be explained exclusively by meals within d_1. Thus, large/small values of $t_0(\boldsymbol{y_k})$ imply that a meal within d_1 is likely/unlikely. Similarly, large/small values of $t_1(\boldsymbol{y_k})$ indicate that a meal within d_0 is likely/unlikely.

In order for the test statistic in Equation 2 to be nontrivial, necessitates the selection of d_0, d_1, δ, and w in Figure 1 such that all dimensions within in $\boldsymbol{G_0}$ and $\boldsymbol{G_1}$ are non-negative. As stated in the text, this work selects d_0, d_1, and δ to be five time steps and w to be 300 time steps. In general, the performance of the PAIN-based detector varies with the selected parameters. Qualitatively, increasing w improves the detector performance as long as the Bergman model remains accurate. At the same time, decreasing d_0, d_1, and δ improves detection accuracy (and time-to-detection) so long as the statistics are nontrivial. Quantifying the detector performance is a subject of future work.