Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Simple Strategies for Retrospective Detection of Meals in Diabetes Datasets

E. Mejia Gamarra [*,***] **F. Reiterer** [*] **P. Tkachenko** [*]
P. Schrangl [*] **G. Freckmann** [**] **W. Ipanaqué** [***]

[*] *Institute for Design and Control of Mechatronical Systems, Johannes Kepler University Linz, Austria (e-mail: pavlo.tkachenko@jku.at).*
[**] *Institut für Diabetes-Technologie Forschungs- und Entwicklungsgesellschaft mbH an der Universität Ulm, Germany*
[***] *Universidad de Piura, Department of Electronics and Automation, Urb. San Eduardo, Piura, Peru*

Abstract: Many model based approaches have been proposed for a personalized insulin therapy in type 1 diabetes (T1D). These approaches rely on patient-specific models of the glucose metabolism which typically need to be identified on high quality data. However, patient data recorded in an at-home setting most often do not meet this criterion, since these are based, among others, on diary entries, which are often erroneous and incomplete. The problem is especially pronounced for recordings of meal intakes which are often accidentally omitted or recorded with wrong time stamps. This paper presents two methods for meal detection based on retrospective analysis of recorded glucose traces. The first method uses the typical signal features of postprandial glucose traces and simple heuristics to detect meals, whereas the second approach relies on similarity measures of glucose traces as compared to postprandial reference profiles. Matching the meal detection results of the algorithms with the actual patient diaries, the methods presented here can be used to find complete, high quality segments in at-home data. Being able to easily distinguish between high and low quality segments in such dataset is expected to improve the reliability of identified patient models.

Keywords: Diabetes, meal detection, biomedical system simulation, system identification

1. INTRODUCTION

Patients with type 1 diabetes (T1D) require regular insulin injections to keep their blood glucose (BG) levels in target, which is needed for reducing the occurrence of diabetes complications. However, hypoglycemia (too low BG values) occurs in case of overdosing insulin, which is unpleasant for the patients and in the worst case can be life threatening. Estimating the required amount of insulin on a day-to-day basis is difficult and a heavy burden for patients with T1D, especially seen that there is a large intrapatient variability of BG dynamics and a myriad of factors that influence the BG level. Therefore, there is a need for tools to assist patients in this task. Great efforts have been invested in the last decades in order to develop algorithms that automatize (parts of) the insulin dosing.

Most of the algorithms proposed in the scientific literature are model-based, meaning that they rely on a model of a patient's glucose metabolism in order to optimize the insulin dosing. Different approaches have been proposed in the literature to obtain a patient-specific model representation, but often it boils down to either optimizing the parameters of a physiological model structure (see *e.g.* Garcia-Tirado et al. (2018)) or to using data-based approaches to identify a black-box model (see *e.g.* Cescon et al. (2015)). The data used for model individualization consists typically of glucose traces recorded by continuous glucose monitoring (CGM) systems together with information about meals (timing and estimated carbohydrate content) and insulin injections (timing and insulin amount). The quality of the used datasets is key to obtain reliable patient representations. This is problematic since data recorded by patients during their every-day life is usually low quality. This problem holds in particular for data which are only available from patient diaries, *i.e.* the meal and insulin recordings. Whereas the use of insulin pumps (which are increasingly widespread among T1D patients) eliminates the need of manually logging insulin injections, meal intakes are only available from diary entries. Patient diaries, however, are often incomplete and erroneous, seen that patients forget to put data into their diary or record meals with a wrong time stamp or inaccurate estimates for the carbohydrate content.

In order to obtain reliable model representations of patients, it seems straightforward that only data segments that are deemed plausible should be used for identification. An algorithm that automatically preselects such segments could be a key element to make personalized model-based insulin therapy approaches better applicable for real-life situations. As a first step into this direction, the current paper introduces simple approaches for the retrospective detection of meals in recorded diabetes datasets.

Meal detection algorithms (MDAs) for diabetes data have already been studied extensively. Most literature on the topic, however, deals with online detection of meals as

a module for artificial pancreas systems, see *e.g.* Dassau et al. (2008); Lee and Bequette (2009); Harvey et al. (2014); Turksoy et al. (2016); Weimer et al. (2016); Samadi et al. (2017, 2018); Kölle et al. (2017); Mahmoudi et al. (2017, 2019); Ramkissoon et al. (2018); Zhao and Zhao (2019). There has hardly been any work on checking meal entries in diabetes datasets for plausibility and completeness. The only work in this context that we are aware of this is Estrada et al. (2009), which, however, has never been validated on gold standard data. Hence, the applicability of the approach remains to be analyzed.

The current paper presents two simple methods for the offline detection of meals in diabetes data based on CGM traces. The first method uses simple heuristic rules to detect meals based on filtered CGM data and estimates of its first and second derivative. The second method on the other hand uses postprandial reference profiles and compares those reference profiles with the measured CGM signals. A distance measure is computed to quantify the similarity between the CGM trace and the reference profiles and to detect meals in case the distance is below a certain threshold. The two meal detection approaches are in a next step validated based on very complete high quality data from a clinical trial that are used as gold standard for quantifying their performance. Additionally, in order to facilitate drawing conclusions from the results of the validation calculations, the performance of the two newly proposed offline MDAs is compared to that of an online MDA from the literature, namely the heuristic approach introduced in Harvey et al. (2014).

2. METHODOLOGY

2.1 Method 1: Simple Heuristic Approach

In order to determine the points in time with meal intakes from CGM traces the following algorithm is used:

(1) Filter the original CGM time signal $y_{\text{cgm}}(t)$ using a Savitzky-Golay-Filter (SGF) (Schafer, 2011). A filter of degree d and with a window width w is used for this purpose. This results in a filtered signal $y(t)$.
(2) Detect all minima and maxima in the filtered signal y.
(3) For each rising segment (minimum until next maximum) the following points in time are computed (see Fig. 1):
   - $t_{\min}$: time where a local minimum occurs in $y(t)$
   - $t_{\max}$: time where the next local maximum occurs in $y(t)$ after $t_{\min}$
   - t_1: time where $\dot{y}(t)$ has largest value in the interval $[t_{\min}, t_{\max}]$
   - t_2: time where $\ddot{y}(t)$ has largest value in the interval $[t_{\min}, t_{\max}]$
   - $\Delta y = y(t_{\max}) - y(t_{\min})$

   If the condition
   $$\Delta y > \Delta y_{\min} \wedge \dot{y}(t_1) > \dot{y}_{\min}$$
   holds (with $\dot{y}_{\min}$ the CGM rate-of-change (ROC) threshold and $\Delta y_{\min}$ the minimum glucose hub), a meal detection is considered at the point in time:
   $$t_{\text{meal}} = 0.5 \cdot (t_{\min} + t_2).$$

The values of all parameter of the algorithm can be found in Tab. 1. It should be mentioned that all MDAs described in the following two subsection are tuned in a way that the specificity of the algorithm is maximized while keeping the overall sensitivity of the algorithms at a comparable level.

Table 1. Parameters of the heuristic offline algorithm for meal detection.

Parameter	Symbol	Unit	Value
SGF degree	d	1	3
SGF window length	w	min	125
ROC threshold	$\dot{y}_{\min}$	mg/dl/min	0.8
Min. glucose hub	$\Delta y_{\min}$	mg/dl	20

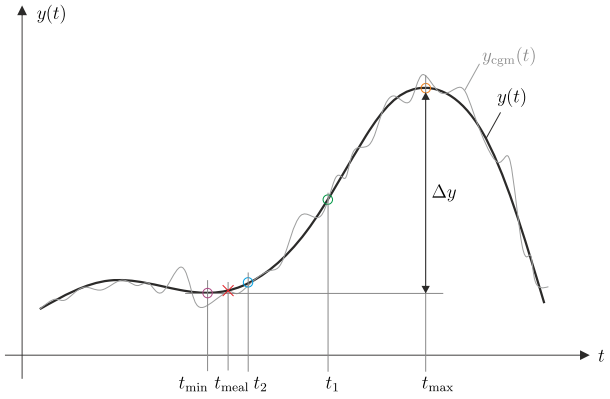

Fig. 1. Illustration of meal detections via the heuristic offline algorithm.

2.2 Method 2: Pattern Matching under Dynamic Time Warping

The idea behind the second method is an (elastic) pattern matching (PM) approach. Given a reference description in terms of in some sense "average" behavior of a CGM trace after the meal (the pattern or template) and sensor data (the input sequence), the goal is to correctly classify the observed data using the catalogue of templates. As the postprandial glucose profile can vary vastly between meals depending on the mixed meal composition (glycemic index, fat and protein content, etc.), computing similarity measures based on the Euclidian distance between reference trace and CGM trace (like *e.g.* the cross-correlation) is expected to result in suboptimal outcomes.

In order to overcome these limitations, the so-called dynamic time warping (DTW) (Wang et al., 2018) of the two sequences can be used, a non-linear alignment technique allowing to set up a similarity measure robust with respect to scaling. DTW, introduced in Sakoe and Chiba (1971), is used for measuring similarity between time series which may vary (*i.e.* warp) in timing.

To describe the DTW method, assume we are provided with two discrete signal sequences, X and Y, of length N and M respectively, where
$$X = [x_1, x_2, \ldots, x_n, \ldots, x_N]$$
$$Y = [y_1, y_2, \ldots, y_m, \ldots, y_M].$$
In order to align two sequences using DTW, one constructs a N-by-M matrix, in which the (n-th, m-th) element is the distance $d(x_n, y_m)$ between the points x_n and y_m (e.g., the Manhattan distance $d(x_n, y_m) = |x_n - y_m|$). A warping

path W is a continuous set of matrix elements that defines a mapping between X and Y. The i-th element of W, $w_i = (n, m)_i$, is defined as a pair of indexes of elements x_n, y_m for which $d(x_n, y_m)$ is computed. As a result,

$$W = [w_1, \ldots, w_i, \ldots, w_I], \quad \max(N, M) \leq I < N + M - 1.$$

Any warping path W should satisfy the following conditions (see for example Kruskal and Liberman (1983)):

(i) **Boundary conditions**: $w_1 = (1, 1)$ and $w_I = (N, M)$. This means that the warping path starts and finishes in diagonally opposite corners of the matrix.
(ii) **Continuity**: Given $w_i = (n, m)_i$, then $w_{i-1} = (n^*, m^*)_{i-1}$, where $n - n^* \leq 1$ and $m - m^* \leq 1$. This restricts the allowable steps in the warping path to adjacent cells (including diagonally adjacent cells).
(iii) **Monotonicity**: Given $w_i = (n, m)_i$, then $w_{i-1} = (n^*, m^*)_{i-1}$, where $n - n^* \geq 0, m - m^* \geq 0$. This forces points in W to be monotonically spaced in time.

Clearly, there are exponentially many warping paths satisfying the above conditions. However, the optimal path is the path minimizing the warping cost:

$$J_{\mathrm{DTW}}(X, Y) = \min \left(\sum_{i=1}^{I} d(x_n, y_m)_i \right), \; i : w_i \in W.$$

This optimal path can be found using dynamic programming to evaluate the following recurrence, which defines the cumulative distance $\gamma(n, m)$ as distance $d(x_n, y_m)$ and the minimum of cumulative distances of adjacent elements:

$$\gamma(n, m) = d(x_n, y_m) + \min\{\gamma(n-1, m-1), \gamma(n-1, m), \gamma(n, m-1)\}.$$

Another crucial issue of the PM algorithm is the definition of a pattern itself. The idea is to use the following strategy for getting a CGM pattern of a meal: First, one extracts a training set of CGM traces after meals. Note that the training set members are not restricted to be of the same length. It is also recommended to normalize the training set patterns – in our case we normalized them to the 0-1 range. Finally, the averaging of the training set elements to get a pattern should be performed. Mathematically, the average (or the mean) $\bar{o}$ of a set of objects O embedded in a space induced by a distance d is

$$\bar{o}(O) = \arg\min_{\bar{o}} \sum_{o \in O} d(\bar{o}, o).$$

If d is the Euclidean distance, the arithmetic mean solves the problem exactly, but for the DTW distance one needs to perform the simultaneous alignment of all training set members. The latter one requires $O(N^L)$ operations, where N is the sequence length and L is the number of sequences. Of course, such complexity is not acceptable.

In order to overcome this difficulty, we use the DTW Barycenter Averaging (DBA) technique from Petitjean et al. (2011). It has the advantages that it avoids using iterative pairwise averaging, does not depend on the order of points and does not increase the length of the resulting pattern. Since the meal composition is known to vary significantly between breakfast, lunch and dinner, and as a result also postprandial glucose profiles differ, separate reference profiles are used for those three meal types. In order to extract meaningful reference profiles, the first two days of CGM data from a clinical trial (Zschornack et al. (2013)) are analyzed. In order to facilitate the comparison of glucose profiles, all CGM traces are smoothed using the SGF described in Sec. 2.1. The extracted reference profiles always start at the minimum in the postprandial glucose (typically exactly at mealtime) and end f minutes after the maximum in the postprandial glucose (breakfast: $f =$ 30 min, lunch: $f = 20$ min, dinner: $f = 120$ min). Based on pre-calculations it was found that it is advantageous to use a different value of f for the three different meal types. The patterns obtained by DBA for breakfasts, lunches and dinners together with the corresponding normalized training sets are illustrated in Fig. 2.

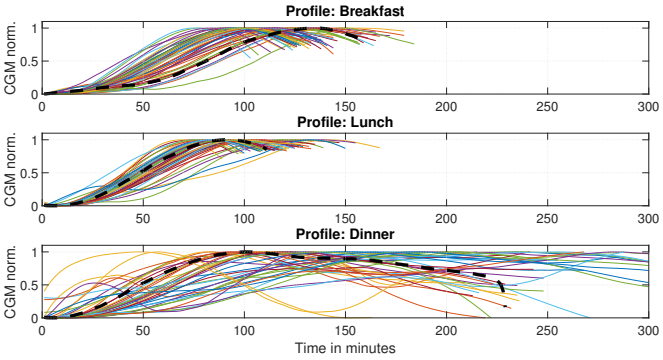

Fig. 2. Illustration of reference profiles used for the meal detection via the DTW approach (reference profiles obtained via DBA are marked as dashed black lines).

The basic idea of PM implementation is the sliding window strategy: assuming sampled CGM measurements, at each time point t_v a window of size v is spanned in the area $[t_v, t_v + v]$. The measured points in the window are used as a test sequence to be compared with a pattern using DTW to make a decision at the point of interest. Based on the value of t_v either the breakfast, lunch or dinner profile from Fig. 2 is used as a reference. A meal is detected when the computed DTW distance is below a certain fixed threshold r. Among all connected CGM points below the threshold, the time of meal detection is placed at the time with minimum DTW distance. A compilation of all settings of the DTW approach for breakfast, lunch and dinner can be found in Tab. 2. An illustrative example for the detection of a breakfast is depicted in Fig. 3. It can be seen that the DTW approach detects the meal event relatively late in this example, with the detection already close to the postprandial peak. This behavior can easily be explained seen that the length of the sliding window v is significantly shorter for breakfast than the length of the reference profile (see Tab. 2 vs. Fig. 2). These settings were found to result in better outcomes regarding sensitivity and specificity of the MDA, whereas the accuracy of the timing of the detections has not explicitly been considered in the current work while tuning the algorithm.

2.3 Online Approach by Harvey

Besides the two newly proposed offline MDAs also an online MDA from the literature is considered in this work for easier comparison, which facilitates rating the performance of the newly proposed algorithms. The Glucose

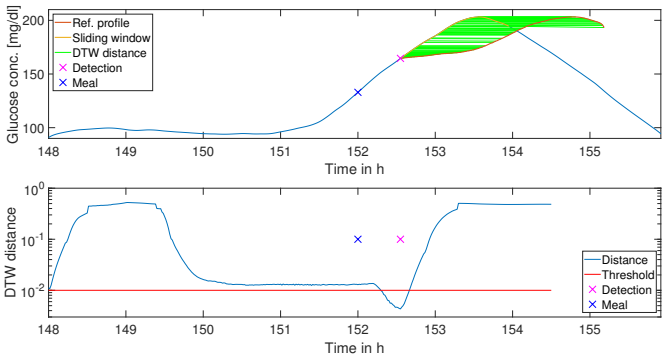

Fig. 3. Example for the DTW approach for offline detection of meals.

Table 2. Parameters of the DTW approach for offline detection of meals.

Parameter	Breakfast	Lunch	Dinner
Starting point of phase	4 a.m	10:30 a.m.	3 p.m.
End point of phase	10:30 a.m.	3 p.m.	4 a.m.
End of ref. profile (max.+f)	30 min	20 min	120 min
Sliding window length, v	85 min	75 min	240 min
DTW distance threshold, r	0.01	0.045	0.02

Rate Increase Detector (GRID) algorithm described in Harvey et al. (2014) is used for this purpose. The algorithm by Harvey is a heuristic approach and detects meals based on estimates of the first derivative of CGM traces. In order to minimize the detrimental effect of measurement errors and sensor noise a noise spike and low pass filter is used to obtain a filtered glucose value G_{F} from which the first derivative G'_{F} is computed. A meal is detected in case the following criteria are fulfilled:

$$\mathrm{MDA}_{\mathrm{Harvey}} = \begin{cases} 1 & \text{if } (G'_{\mathrm{F}}(k{-}2{:}k) > G'_{\mathrm{min},3}) \\ & \vee\ (G'_{\mathrm{F}}(k{-}1{:}k) > G'_{\mathrm{min},2}) \\ 0 & \text{otherwise} \end{cases}$$

with threshold values $G'_{\mathrm{min},2}$ and $G'_{\mathrm{min},3}$. These thresholds are set to $G'_{\mathrm{min},2} = 1.9$ mg/dl/min and $G'_{\mathrm{min},3} = 1.7$ mg/dl/min. An additional criterion is introduced in the current work that prevents the algorithm to switch to 1 in case there has been a previous meal detection within the last 120 minutes, thereby reducing the number of false positive meal detections.

3. RESULTS AND DISCUSSION

3.1 Clinical Data

For the current work data from a clinical trial performed at the Institute for Diabetes Technology, Germany is used (Zschornack et al. (2013)). In this clinical trial 37 subjects with T1D spent seven days hospitalized. During this time period each of them wore either two (28 individuals) or four (9 individuals) CGM systems in parallel. During the entire period of the study all CGM signals have been recorded, together with frequent BG measurements by means of self-monitoring of blood glucose (SMBG, at least one measurement per hour during the day) and detailed documentation about meal intakes (together with the corresponding carbohydrate content of the meals), bolus insulin injections and basal insulin rates. Information about meal timing, size and composition (among others: carbohydrate content) are highly reliable seen that they have been recorded by the study staff and ingested meals have been analyzed by a trained dietitian.

3.2 Performance evaluation

For the performance assessment the following indicators are considered:

- Sensitivity (SE): Defined as the percentage of meals that are detected by the algorithm. A meal is considered detected if the MDA has a detection event within 75 minutes around a meal ingestion. If this is not the case this counts as a false negative (FN).
- FP/day: Average number of false positive (FP) meal detections per day. A detection event counts as FP if there is no meal within 75 minutes around the detection.
- ΔT: Average absolute time between a meal ingestion and the detection event.

As the by Harvey is an online algorithm and has a different mode of operation than the proposed ones, different criteria have to be used to distinguish between true positive (TP), FP and FN detection events. A TP detection is counted if there is a meal detection within 120 min of a meal ingestion, whereas a FP event is defined as a meal detection without any ingested meal within 120 min prior to the detection event.

For all studied algorithms a special treatment is performed regarding meal detections of small meals with a carbohydrate content of below 20 g. These snacks do not contribute to FN detections, *i.e.* in case a snack does not trigger a meal flag by the MDAs this does not count as a FN. On the other hand, in case there is a meal detection after a snack, this is indeed counted as a TP event.

An illustrative result for the performance of the different MDAs is plotted in Fig. 4. Meals are marked with (+) in this plot, whereas detection events are marked with (×). The green vertical lines correspond to midnight. In general, all approaches work relatively well for the selected data segment, especially for the breakfast events. As can be expected the online method by Harvey results in meal detection that are always after the meal ingestion, whereas for the offline methods detections can be both before or after meal meals. For breakfasts, however, it can be seen that the DTW approach is always significantly late with its detections. As already discussed earlier, this behavior is due to the difference between the sliding window length v and the length of reference profile.

Besides computing the overall average performance measures for all 37 patients, additionally, results are also analyzed for breakfast, lunch and dinner separately. Breakfasts are defined as all meal events that are reported between 4 a.m. and 10:30 a.m., lunch as meals that are recorded between 10:30 a.m. and 3 p.m., whereas all meals recorded between 3 p.m. and 4 a.m. are classified as dinner events. The results of this analysis are reported in Tab. 3 as average ± standard deviation (average over the mean values per patient).

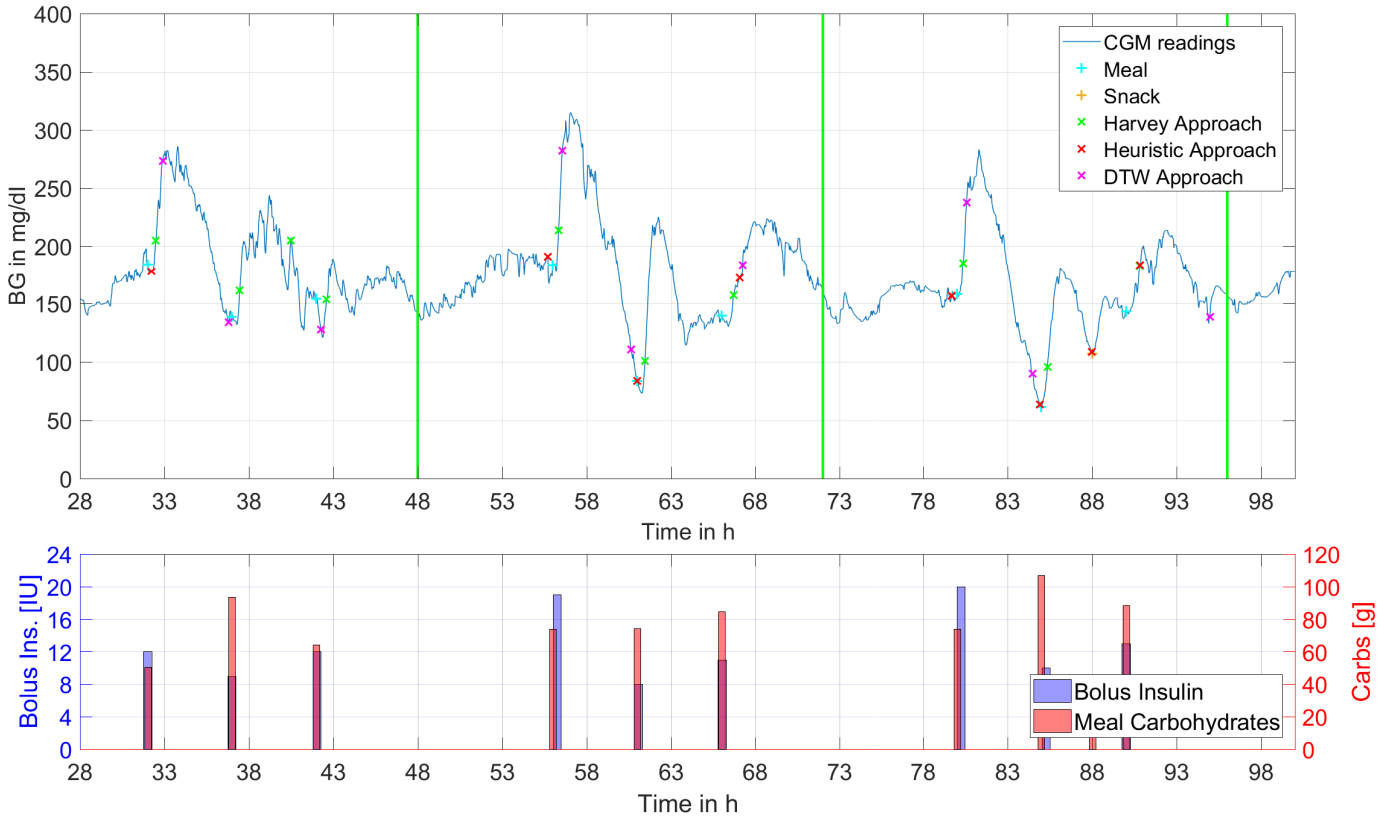

Fig. 4. Illustrative results regarding performance of studied meal detection algorithms (patient P0012).

From the results in Tab. 3 it can be seen that the absolute time difference between meal detection and meal event is the smallest for the newly proposed heuristic approach. For the DTW method the time differences are even slightly bigger than for the online approach by Harvey. If SE and specificity of the algorithms are compared based only the rows with the overall results in Tab. 3, it may appear as if the advantages of the newly proposed offline MDAs are not very big as compared to the online algorithm by Harvey. Nevertheless, already from the overall results it becomes apparent that the newly proposed heuristic approach does lead to a significantly lower number of FP detections compared to the Harvey algorithm. For the DTW approach on the other hand no significant advantages can be seen in the overall results, but the strengths of the approach become visible if breakfast results are analyzed separately: For breakfast events the DTW approach has by far the smallest number of FP detections and the highest SE, whereas differences between the heuristic offline approach and the method by Harvey are relatively small. For lunch events all analyzed algorithms show comparably good results. For dinner, however, none of the algorithms show a satisfactory level of SE. This can easily be explained by the fact that postprandial profiles are more diverse for dinner events (see also Fig. 2), making the meal detection task the most difficult. Regarding the number of FP detections at dinnertime, the heuristic offline approach shows the best results among the analyzed algorithms.

Overall, it can be stated that the two newly proposed offline MDAs already give promising results. Further work should be put though into increasing the detection sensitivity of algorithms, especially around dinnertime, but without any further increase in FP events. For the DTW approach it should furthermore be investigated how meal timing can be estimated more accurately.

Table 3. Comparison of overall MDA results.

		ΔT [min]	FP/day [-]	SE [%]
Heuristics	Breakfast	21.7 ± 9.0	0.15 ± 0.17	78.5 ± 23
	Lunch	16.2 ± 7.1	0.04 ± 0.11	83.1 ± 19
	Dinner	22.6 ± 11.5	0.23 ± 0.31	61.1 ± 20
	Overall	19.5 ± 4.3	0.42 ± 0.43	72.5 ± 13
DTW	Breakfast	27.8 ± 9.4	0.07 ± 0.13	90.0 ± 15
	Lunch	32.0 ± 10.3	0.05 ± 0.12	85.0 ± 16
	Dinner	43.1 ± 13.6	0.42 ± 0.38	59.3 ± 22
	Overall	34.6 ± 6.4	0.55 ± 0.45	75.2 ± 13
Harvey	Breakfast	27.8 ± 8.7	0.18 ± 0.22	80.6 ± 20
	Lunch	28.4 ± 6.4	0.06 ± 0.11	80.0 ± 14
	Dinner	39.4 ± 10.6	0.34 ± 0.37	62.3 ± 21
	Overall	31.0 ± 5.2	0.54 ± 0.52	73.3 ± 12

4. CONCLUSIONS AND OUTLOOK

The current paper presents two new algorithms for the offline detection of meals in diabetes data and validates them based on gold standard data from a clinical trial. This work is a first step towards an algorithm assessing the quality of recorded outpatient diabetes datasets. Such an algorithm would facilitate the usage of low quality at-home data for the identification of personalized models of glucose

metabolism and thus could be key for the applicability of many personalized insulin therapy approaches.

One of the further steps to be taken is the automatized estimation of the carbohydrate content of meals. In diabetes datasets possible flaws are not only missing meals and meals with incorrect timestamps, but there might also be meals which are recorded with grossly underestimated or overestimated meal size. Such errors can currently not be detected with the newly introduced algorithms, since these only estimate the timing of meal events, but do not give information about meal size. A corresponding extension of the offline MDAs is envisioned in future works.

A further step would also include an analysis of diabetes data with respect to bolus insulin injections. The two algorithms presented in the current work do not consider any information about bolus injections. In case the recorded bolus intakes can be regarded as certain (*i.e.* if those are logged automatically by either an insulin pump or a smart pen), the task of retrospective meal detection becomes much easier seen that meals and bolus injections are typically highly correlated in time, but also in size. In the common case when patients use standard insulin pens the information about injections still needs to be recorded manually and the corresponding data quality can be expected to be similarly low as for meal intakes.

Finally, the basic hypothesis behind the current work remains to be verified. Even though it seems straightforward that model quality is improved if only a subset of the available outpatient data with plausible input information is used for personalization, this remains to be demonstrated.

REFERENCES

Cescon, M., Johansson, R., and Renard, E. (2015). Subspace-based linear multi-step predictors in type 1 diabetes mellitus. *Biomedical Signal Processing and Control*, 22, 99–110.

Dassau, E., Bequette, B.W., Buckingham, B.A., and Doyle, F.J. (2008). Detection of a meal using continuous glucose monitoring (cgm): implications for an artificial β-cell. *Diabetes Care*.

Estrada, G.C., Kirchsteiger, H., Del Re, L., and Renard, E. (2009). Model based validation of meal inputs in diabetes therapy. In *Proceedings of the 15th IFAC Symposium on System Identification (SYSID)*, 239–244.

Garcia-Tirado, J., Zuluaga-Bedoya, C., and Breton, M.D. (2018). Identifiability analysis of three control-oriented models for use in artificial pancreas systems. *Journal of Diabetes Science and Technology*, 12(5), 937–952.

Harvey, R.A., Dassau, E., Zisser, H., Seborg, D.E., and Doyle III, F.J. (2014). Design of the glucose rate increase detector: a meal detection module for the health monitoring system. *Journal of Diabetes Science and Technology*, 8(2), 307–320.

Kölle, K., Fougner, A.L., and Stavdahl, Ø. (2017). Meal detection based on non-individualized moving horizon estimation and classification. In *2017 IEEE Conference on Control Technology and Applications (CCTA)*, 529–535.

Kruskal, J. and Liberman, M. (1983). The symmetric time-warping problem: From continuous to discrete.

Lee, H. and Bequette, B.W. (2009). A closed-loop artificial pancreas based on model predictive control: Human-friendly identification and automatic meal disturbance rejection. *Biomedical Signal Processing and Control*, 4(4), 347–354.

Mahmoudi, Z., Cameron, F., Poulsen, N.K., Madsen, H., Bequette, B.W., and Jørgensen, J.B. (2019). Sensor-based detection and estimation of meal carbohydrates for people with diabetes. *Biomedical Signal Processing and Control*, 48, 12–25.

Mahmoudi, Z., Nørgaard, K., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2017). Fault and meal detection by redundant continuous glucose monitors and the unscented kalman filter. *Biomedical Signal Processing and Control*, 38, 86–99.

Petitjean, F., Ketterlin, A., and Gancarski, P. (2011). A global averaging method for dynamic time warping, with applications to clustering. *Pattern Recognition*, 44(3), 678 – 693.

Ramkissoon, C.M., Herrero, P., Bondia, J., and Vehi, J. (2018). Unannounced meals in the artificial pancreas: Detection using continuous glucose monitoring. *Sensors*, 18(3), 884.

Sakoe, H. and Chiba, S. (1971). A dynamic programming approach to continuous speech recognition. In *Proceedings of the Seventh International Congress on Acoustics, Budapest*, volume 3, 65–69. Akadémiai Kiadó, Budapest.

Samadi, S., Rashid, M., Turksoy, K., Feng, J., Hajizadeh, I., Hobbs, N., Lazaro, C., Sevil, M., Littlejohn, E., and Cinar, A. (2018). Automatic detection and estimation of unannounced meals for multivariable artificial pancreas system. *Diabetes Technology & Therapeutics*, 20(3), 235–246.

Samadi, S., Turksoy, K., Hajizadeh, I., Feng, J., Sevil, M., and Cinar, A. (2017). Meal detection and carbohydrate estimation using continuous glucose sensor data. *IEEE Journal of Biomedical and Health Informatics*, 21(3), 619–627.

Schafer, R.W. (2011). What is a savitzky-golay filter? [lecture notes]. *IEEE Signal Processing Magazine*, 28(4), 111–117.

Turksoy, K., Samadi, S., Feng, J., Littlejohn, E., Quinn, L., and Cinar, A. (2016). Meal detection in patients with type 1 diabetes: a new module for the multivariable adaptive artificial pancreas control system. *IEEE Journal of Biomedical and Health Informatics*, 20(1), 47–54.

Wang, W., Lyu, G., Shi, Y., and Liang, X. (2018). Time series clustering based on dynamic time warping. *2018 IEEE 9th International Conference on Software Engineering and Service Science (ICSESS)*, 487–490.

Weimer, J., Chen, S., Peleckis, A., Rickels, M.R., and Lee, I. (2016). Physiology-invariant meal detection for type 1 diabetes. *Diabetes Technology & Therapeutics*, 18(10), 616–624.

Zhao, H. and Zhao, C. (2019). A concurrent fault and meal detection method based on dynamics analysis for continuous glucose monitoring sensor. *Chemometrics and Intelligent Laboratory Systems*, 189, 72–80.

Zschornack, E., Schmid, C., Pleus, S., Link, M., Klötzer, H.M., Obermaier, K., Schoemaker, M., Strasser, M., Frisch, G., Schmelzeisen-Redeker, G., Haug, C., and Freckmann, G. (2013). Evaluation of the performance of a novel system for continuous glucose monitoring. *Journal of Diabetes Science and Technology*, 7(4), 815–823.