

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2017206786 B2**

(54) Title
Predictive control model for the artificial pancreas using past predictions

(51) International Patent Classification(s)
A61M 5/142 (2006.01) **G01N 33/53** (2006.01)
A61P 3/10 (2006.01) **G01N 33/66** (2006.01)
G01N 33/50 (2006.01)

(21) Application No: **2017206786** (22) Date of Filing: **2017.01.12**

(87) WIPO No: **WO17/123805**

(30) Priority Data

(31) Number	(32) Date	(33) Country
62/277,788	2016.01.12	US

(43) Publication Date: **2017.07.20**

(44) Accepted Journal Date: **2022.01.27**

(71) Applicant(s)
President and Fellows of Harvard College;The Regents of the University of California

(72) Inventor(s)
Laguna Sanz, Alejandro J.;Dassau, Eyal;Doyle III, Francis J.

(74) Agent / Attorney
FB Rice Pty Ltd, Level 23 44 Market Street, Sydney, NSW, 2000, AU

(56) Related Art
US 2015/0100038 A1
US 2014/0276555 A1
US 2007/0244575 A1



(51) International Patent Classification:

A61M 5/142 (2006.01) G01N 33/53 (2006.01)
A61P 3/10 (2006.01) G01N 33/66 (2006.01)
G01N 33/50 (2006.01)

President and Fellows of Harvard College, 17 Quincy Street, Cambridge, Massachusetts 02138 (US). **DOYLE, Francis, J., III**; c/o President and Fellows of Harvard College, 17 Quincy Street, Cambridge, Massachusetts 02138 (US).

(21) International Application Number:

PCT/US2017/013248

(74) Agents: **RESNICK, David, S.** et al.; Nixon Peabody LLP, 100 Summer St., Boston, Massachusetts 02110-2131 (US).

(22) International Filing Date:

12 January 2017 (12.01.2017)

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/277,788 12 January 2016 (12.01.2016) US

(71) Applicants: **PRESIDENT AND FELLOWS OF HARVARD COLLEGE** [US/US]; 17 Quincy Street, Cambridge, Massachusetts 02138 (US). **THE REGENTS OF THE UNIVERSITY OF CALIFORNIA** [US/US]; 1111 Franklin Street, Oakland, California 94607 (US).


(72) Inventors: **LAGUNA SANZ, Alejandro**; c/o President and Fellows of Harvard College, 17 Quincy Street, Cambridge, Massachusetts 02138 (US). **DASSAU, Eyal**; c/o

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ,

[Continued on next page]

(54) Title: PREDICTIVE CONTROL MODEL FOR THE ARTIFICIAL PANCREAS USING PAST PREDICTIONS

FIG. 8



(57) Abstract: Presented herein are methods, and devices of adaptively tuning a zone based Model Predictive Control (MPC) controller, using at least one processor, which include determining, residuals based on prediction models storing, in a memory, the determined residuals, calculating a trust index by quantifying uncertainty of the prediction models using the stored residuals and tuning the MPC controller, in real time based on the calculated value of the trust index.

WO 2017/123805 A1

TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

PREDICTIVE CONTROL MODEL FOR THE ARTIFICIAL PANCREAS USING PAST PREDICTIONS

Background***Technical Field***

[0001] The present disclosure relates generally to model predictions of glucose concentration for artificial pancreas. More particularly, aspects of this disclosure relate to methods, and devices directed towards an adaptive model predictive controller to improve model prediction accuracy.

Description of Related Art

[0002] The destruction of the insulin secretion pancreatic beta-cells in people with type 1 diabetes mellitus causes glucose concentration to increase to unhealthy threatening levels in the bloodstream. Prolonged high blood glucose concentration is harmful to the patient, making an exogenous insulin delivery a necessity. However, excessive insulin delivery is also dangerous, since too low glucose concentration (hypoglycemia) can quickly lead to coma or even death. The most common treatment historically has been Multiple Daily Injections (MDI) of insulin to correct the ever-increasing glucose concentration. Even though MDI is still the most applied treatment, the introduction of Continuous Subcutaneous Insulin Infusion (CSII) pumps and Continuous Glucose Monitors (CGM) in the last decades has boosted the application of a variety of new, more efficient insulin-based treatments.

[0003] The artificial pancreas is a device that aims to automate insulin delivery (close the loop) by continuously adjusting insulin delivery based on glucose fluctuations. Massive advances have been done in recent years with several research teams completing closed-loop experiments in controlled environments, and more recently in an ambulatory setting and even free living / at home.

[0004] The core of any artificial pancreas is a control algorithm that calculates the appropriate insulin dosing to maintain glucose concentration at a safe healthy level at all times. The most common controllers used in the artificial pancreas are algorithms based on: Model predictive control (MPC), Proportional-Integral-Derivative (PID) approach and Fuzzy-Logic design.

[0005] MPC controllers use prediction models to estimate the trajectories of output variables. However, glucose predictions by mathematical models are still very inaccurate, and model individualization is unsatisfactory for data-based models or physiology-based models,

resulting in great uncertainty in the model predictions for individual patients. Several recent studies about the quantification of the uncertainty have been published by bounding the uncertainty using interval identification, showing very promising results.

[0006] MPC performance depends on the accuracy of the prediction models implemented for the controller. Prediction errors must be measured and accounted for in the glucose regulation process. Thus, there is a need for an MPC algorithm that is personalized to the patient's clinical information and an MPC with a reference defined as a glucose zone, designed to deliver the optimal insulin dose that brings glucose predictions to a healthy range within a limited time frame.

Summary

[0007] According to an aspect of an exemplary embodiment, a method of adaptively tuning a zone based Model Predictive Control (MPC) controller, using at least one processor, includes determining, using at least one of said at least one processor, residuals based on prediction models, storing, in a memory, the determined residuals, calculating a trust index, using at least one of said at least one processor, by quantifying uncertainty of the prediction models using the stored residuals and tuning the MPC controller, in real time, using at least one of said at least one processor, based on the calculated value of the trust index.

[0008] According to another exemplary embodiment the determining further includes determining the residual at time t and for a prediction k steps ahead using the formula:

$$res_k^t = y_{k|t-k} - CGM_t \quad \forall k \in \mathbb{Z}_1^{N_p}$$

where $y_{k|t-k}$ represents predicted glucose by the model at time t that was predicted k samples in the past, N_p represents prediction horizon and CGM_t represents current glucose values.

[0009] According to another exemplary embodiment the storing further comprises storing each of the N_p residuals (res_k^t) in a pool of most recent N_f residuals, resulting in a matrix $N_p \times N_f$.

[0010] According to another exemplary embodiment, the method further includes penalizing, using at least one of said at least one processor, the determined residuals using a forgetting function, where the forgetting function is determined using the formula:

$$\tilde{res}_{k,i}^t = res_k^{t-i} \cdot \frac{\log(N_f - i + 1)}{\log N_f} \quad \forall i \in \mathbb{Z}_1^{N_f}, \forall k \in \mathbb{Z}_1^{N_p}$$

where t represents current time, res_k^{t-i} represent progressively more distant residuals, and $\widetilde{res}_{k,i}^t$ represents the penalized residual.

[0011] According to another exemplary embodiment, for each timestamp t each residual is penalized differently.

[0012] According to another exemplary embodiment, the tuning further comprises tuning the MPC conservatively when the determined residuals have a high value.

[0013] According to another exemplary embodiment, the tuning further comprises tuning the MPC conservatively when the penalized residuals have a high value.

[0014] According to another exemplary embodiment, the tuning further comprises tuning the MPC aggressively when the determined residuals have a low value.

[0015] According to another exemplary embodiment, the tuning further comprises tuning the MPC aggressively when the penalized residuals have a low value.

[0016] According to another exemplary embodiment, the method further includes calculating, using at least one of said at least one processor, an estimate of a confidence interval of a current prediction error based on the penalized residuals using the formula:

$$B_x = \max(\text{percentile}_x(\widetilde{res}_k^t), 0) \quad \forall k \in \mathbb{Z}_1^{N_p}$$

where B_k is an empirically defined boundary of the confidence interval of a prediction error k steps ahead, and $\text{percentile}_x(\widetilde{res}_k^t)$ represents x -th percentile function of the penalized residual at time t and for a prediction k steps ahead.

[0017] According to another exemplary embodiment, the calculating the estimate of the confidence interval comprises calculating an empirically defined upper boundary \hat{B}_k and an empirically defined lower boundary \check{B}_k using the formula:

$$\hat{B}_k = \max(\text{percentile}_{95}(\widetilde{res}_k^t), 0) \quad \forall k \in \mathbb{Z}_1^{N_p}$$

$$\check{B}_k = \min(\text{percentile}_5(\widetilde{res}_k^t), 0) \quad \forall k \in \mathbb{Z}_1^{N_p}$$

[0018] According to another exemplary embodiment, the calculating the trust index comprises calculating the trust index based on the calculated empirically defined upper boundary \hat{B}_k and the calculated empirically defined lower boundary \check{B}_k using the formula:

$$T_t := \hat{B}_k - \check{B}_k + d_H(\text{CGM}_t, [\check{B} + y_t, \hat{B} + y_t])$$

$$d_H(x, [a, b]) := \begin{cases} 0 & \text{if } x \in [a, b] \\ x - b & \text{if } x > b \\ a - x & \text{if } x < a \end{cases}$$

[0019] where $d_H(x, [a, b])$ represents Hausdorff distance from a point x to an interval $[a, b]$, and y_t represents model's prediction of current CGM sample.

[0020] According to another aspect of an exemplary embodiment, a method of insulin delivery using an artificial pancreas, the artificial pancreas including at least one processor, includes determining, using at least one of said at least one processor, residuals based on prediction models, storing, in a memory, the determined residuals, calculating a trust index, using at least one of said at least one processor, by quantifying uncertainty of the prediction models using the stored residuals, tuning a model prediction control (MPC) controller, in real time, using at least one of said at least one processor, based on the calculated value of the trust index, and controlling, using at least one of said at least one processor, the delivery of insulin based on the tuned MPC controller.

[0021] According to another aspect of an exemplary embodiment, an artificial pancreas apparatus for insulin delivery includes at least one non-transitory memory operable to store program code, a model prediction control (MPC) controller including at least one processor operable to read said program code and operate as instructed by said program code, said program code causing the at least one processor to determine residuals based on prediction models, store, in the at least one non-transitory memory, the determined residuals, calculate a trust index by quantifying uncertainty of the prediction models using the stored residuals, tune the MPC controller in real time based on the calculated value of the trust index, and deliver the insulin based on the tuned MPC controller.

[0021A] According to another aspect of an exemplary embodiment, a method of adaptively tuning a Model Predictive Control (MPC) controller, using at least one processor, the method comprising: determining, using at least one of said at least one processor, a plurality of residuals, wherein each of the plurality of residuals is based on a difference between a measured glucose value and a glucose value predicted by a model at a previous time among of a plurality of time samples in the past; storing, in a memory, the determined residuals; calculating a trust index, using at least one of said at least one processor, by quantifying uncertainty of the model using the stored residuals; tuning the MPC controller, in real time, using at least one of said at least one processor, based on the calculated value of the trust index; and determining an amount of insulin to be delivered by an insulin pump based on insulin amounts determined by the tuned MPC controller.

[0021B] According to another aspect of an exemplary embodiment, a method of insulin delivery using an artificial pancreas, the artificial pancreas including at least one processor, the method comprising: determining, using at least one of said at least one processor, a plurality of

2017206786 29 Jul 2021

residuals, wherein each of the plurality of residuals is based on a difference between a measured glucose value and a glucose value predicted by a model at a previous time among of a plurality of time samples in the past; storing, in a memory, the determined residuals; calculating a trust index, using at least one of said at least one processor, by quantifying uncertainty of the model using the stored residuals; tuning a model prediction control (MPC) controller, in real time, using at least one of said at least one processor, based on the calculated value of the trust index; and controlling, using at least one of said at least one processor, the delivery of insulin based on insulin amounts determined by the tuned MPC controller.

[0021C] According to another aspect of an exemplary embodiment, an artificial pancreas apparatus for insulin delivery, the artificial pancreas comprising: at least one non-transitory memory operable to store program code; a model prediction control (MPC) controller including at least one processor operable to read said program code and operate as instructed by said program code, said program code causing the at least one processor to: determine a plurality of residuals, wherein each of the plurality of residuals is based on a difference between a measured glucose value and a glucose value predicted by a model at a previous time among of a plurality of time samples in the past; store, in the at least one non-transitory memory, the determined residuals; calculate a trust index by quantifying uncertainty of the prediction models using the stored residuals; tune the MPC controller in real time based on the calculated value of the trust index; and determine an amount of insulin to be delivered by an insulin pump of the artificial pancreas based on the tuned MPC controller.

Brief Description of the Drawings

[0022] FIG. 1 depicts a logarithmic forgetting function applied to the residuals, according to an exemplary embodiment.

[0023] FIG. 2 depicts a relationship between confidence on the current measurement and the aggressiveness of the current tuning, according to the adaptive MPC controller.

[0024] FIG. 3 depicts a table describing controller behavior based on the glucose scenario and trust index value, according to an exemplary embodiment.

[0025] FIG. 4 depicts a control variability grid analysis for 10 in silico subjects from the UV A/Padova metabolic simulator using an aggressive tuning of the MPC controller, according to an exemplary embodiment.

[0026] FIG. 5 depicts a control variability grid analysis for 10 in silico subjects from the UV A/Padova metabolic simulator using a conservative tuning of the MPC controller, according to an exemplary embodiment.

[0027] FIG. 6 depicts a control variability grid analysis for 10 in silico subjects from the UVA/Padova metabolic simulator using an adaptive tuning of the MPC controller, according to an exemplary embodiment.

[0028] FIG. 7 depicts a table providing a performance comparison metric of an MPC controller tuned aggressively, conservatively and adaptively, according to an exemplary embodiment.

[0029] FIG. 8 depicts a chart providing a comparison in median glucose concentration and interquartile range (top) and mean insulin pump administration (bottom) between aggressive static tuning (plotted in blue - continuous line) and adaptive tuning (plotted in yellow - discontinuous line) of an MPC controller, the simulated protocol including three unannounced meals compensated using the glucose rate increase detector (GRID) meal detection algorithm, according to an exemplary embodiment.

[0030] FIG. 9 depicts a chart providing a comparison in median glucose concentration and interquartile range (top) and mean insulin pump administration (bottom) between conservative static tuning (plotted in blue - continuous line) and adaptive tuning (plotted in yellow - discontinuous line) of an MPC controller, the simulated protocol including three unannounced meals compensated using the glucose rate increase detector (GRID) meal detection algorithm, according to an exemplary embodiment.

[0031] FIG. 10 depicts a block diagram of an apparatus incorporating an MPC controller, according to an exemplary embodiment.

[0032] FIG. 11 depicts a flowchart describing the process to adaptively tune an MPC controller, according to an exemplary embodiment.

[0033] FIG. 12 depicts adaptation scheme for parameters \hat{R} (top) and \check{R} (bottom).

[0034] FIG. 13 depicts the Mean and standard deviation of the control performance metrics of the zMPC controller for the Default, Aggressive, Conservative, and Adaptive tuning, according to an exemplary embodiment.

[0035] The present disclosure is susceptible to various modifications and alternative forms, and some representative embodiments have been shown by way of example in the drawings and will be described in detail herein. It should be understood, however, that the inventive aspects are not limited to the particular forms illustrated in the drawings. Rather, the disclosure is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the disclosure as defined by the appended claims.

Detailed Description of the Drawings

[0036] Embodiments will be described below in more detail with reference to the accompanying drawings. The following detailed descriptions are provided to assist the reader in gaining a comprehensive understanding of the methods, apparatuses, and/or systems described herein and equivalent modifications thereof. Accordingly, various changes, modifications, and equivalents of the methods, apparatuses, and/or systems described herein will be apparent to those of ordinary skill in the art. Moreover, descriptions of well-known functions and constructions may be omitted for increased clarity and conciseness.

[0037] The terms used in the description are intended to describe embodiments only, and shall by no means be restrictive. Unless clearly used otherwise, expressions in a singular form include a meaning of a plural form. In the present description, an expression such as “comprising” or “including” is intended to designate a characteristic, a number, a step, an operation, an element, a part or combinations thereof, and shall not be construed to preclude any presence or possibility of one or more other characteristics, numbers, steps, operations, elements, parts or combinations thereof.

[0038] It will be understood to those skilled in the art that the disclosure described herein may apply to any type of special-purpose or general-purpose computer, including a standalone computer or portion thereof.

[0039] The present disclosure describes quantifying the model prediction accuracy by a rigorous analysis of the prediction residuals, and calculating a “Trust Index,” to be used as a tuning tool in a closed-loop algorithm. The integration of the information from the prediction residuals into the control algorithm is used to achieve a new generation of adaptive control algorithms. These algorithms act aggressively to eliminate both hyper-glycemia and hypoglycemia in case of good confidence in the glucose profile, determined by a continuous glucose monitor (CGM) according to an exemplary embodiment, and act conservative in case of poor CGM quality, in-turn providing a safer and more effective glucose control.

[0040] The above described results may be achieved by calculating optimal insulin dosage by solving an optimization problem in which a scalar index is minimized by penalizing relative input deviations and glucose predictions out of the reference zone. The controller’s tuning parameters are the penalizations of the input variable (insulin). Positive and negative relative inputs may be penalized differently, as will be described in greater detail below.

[0041] The in silico simulations and results analyzed in the various figures below were achieved using Matlab R2013b (Mathworks, Natick,MA) as the platform. However, other

such programs might be used for the same simulations and analysis. The p-values for statistical comparison were calculated using non-parametric random permutations methods. The datasets were generated using the Universities of Virginia / Padova metabolic simulator, which has been accepted by the FDA in lieu of animal trials. The simulator implements ten parameter combinations that represent the inter-patient variability equivalent to the general diabetic population. A key part of the simulator is a model of the CGM noise based on different probabilistic distributions. The CGM error is different for each execution of the software, introducing an even larger variability in every simulation, even for simulations of the same *in silico* patient. The variability introduced by the random number generator for each simulation impedes a clear comparison of the ten *in silico* patients under different treatments. To minimize this effect, every *in silico* subject is simulated ten times under every scenario to sum a total of 100 datasets. By doing this, the noise introduced by the CGM error model is averaged over the different simulations of each *in silico* subject, while still maintaining the variability that the CGM model introduces on the experiments. The above described scenarios are used as exemplary embodiment to achieve simulations and compare results between statically and adaptively tuned MPC controlled. It should be noted that the disclosure is not limited to the above exemplary embodiment and different datasets, simulators and methods may be used to achieve the simulations and results.

[0042] The simulations, under this exemplary embodiment depicted in the drawings, consist in twenty four hours of CGM data with three unannounced meals of increasing carbohydrates load (50 g, 75 g and 100 g). Every simulation was preceded by two hours of open-loop monitoring as a warm up period. During that time the *in silico* subject received the appropriate basal insulin. In order to minimize the postprandial glucose excursion a meal detection algorithm was used (Glucose Rate Increase Detector, GRID). This algorithm triggered an alarm when certain conditions of glucose concentration and glucose trend were met, and an insulin compensation dose was administered, equivalent to a prandial bolus strategy. The compensation bolus was the equivalent to a 75 g carbohydrate (CHO) meal correction, based on the patient's insulin-to-carbohydrate ratio. GRID also incorporates routines for patient safety including an algorithm that avoids delivering extra insulin according to recent insulin delivery history, and a security interlock in case of missing CGM samples.

[0043] The chosen controller, Zone MPC controller (zMPC), was developed originally by Grosman *et al.* using a constant glucose reference band, and implementing very simple models for the insulin and meal inputs to the system. This control scheme was later expanded

by improving the prediction model, using a control oriented prediction strategy proposed by van Heusden *et al.* Further functionalities were added to the controller and to the prediction model by Gondhalekar *et al.*, implementing different penalization weights for insulin delivery larger or smaller than the patient's basal rate. The cost function for the MPC problem is explained below. The formulae provided below are based on the conventional controllers provided in the above references:

$$J(\cdot) := \sum_{k=1}^{N_p} z_k^2 + \sum_{k=0}^{N_u-1} (\hat{R}\hat{u}_k^2 + \check{R}\check{u}_k^2) \quad 1$$

$$\hat{u}_k := \max(u_k - u_{BASAL}, 0) \quad \forall k \in \mathbb{Z}_0^{N_u-1} \quad 2$$

$$\check{u}_k := \min(u_k - u_{BASAL}, 0) \quad \forall k \in \mathbb{Z}_0^{N_u-1} \quad 3$$

where N_p prediction horizon, and N_u is the control horizon and z_k is the output error, understood as the distance of the glucose prediction to the glucose safe zone. u_k is the insulin input, and u_{BASAL} is the patient's basal insulin.

[0044] The current disclosure describes using \hat{R} and \check{R} parameters to adaptively tune the controller. Both parameters modulate the insulin delivery, but \hat{R} penalizes insulin actions greater than the basal rate of the patient, while \check{R} acts on the insulin action when it is being suppressed below basal rate. Increasing the value of \hat{R} penalizes higher insulin values and tends to make the controller more conservative, delivering less insulin. Increasing the value of \check{R} impedes insulin attenuation / suspension, tending to give more insulin, making the controller more aggressive.

[0045] An aggressive controller is prone to deliver more insulin. It would act fast to counter hyperglycemic excursions, and it would be slow to suspend insulin infusion in presence of hypoglycemia.

[0046] On the other hand, a conservative controller delivers less insulin overall. It shuts down pump infusion fast to minimize hypoglycemia, but it reacts slow to hyperglycemic excursions.

[0047] According to an exemplary embodiment, two different sets of parameters are selected: 1) very aggressive design $\hat{R}=500$, $\check{R}=300$; 2) very conservative design $\hat{R}=50000$, $\check{R}=10$. These will be used to compare the simulation results between statically tuned and adaptively tuned MPC controllers.

[0048] It should be noted that the ratio between \hat{R} and \check{R} provides a measurement of the asymmetry of the controller, but this ratio is not enough to characterize the aggressiveness of the controller. The magnitude of \hat{R} and \check{R} separately penalizes the input costs (as shown in equation 1) relatively to the weight of the output error, which is fixed.

[0049] Referring now to the drawings, wherein like reference numerals refer to like features throughout the several views, there is shown in FIG. 1, a logarithmic forgetting function applied to the residuals, according to an exemplary embodiment.

[0050] The method and device described in the current disclosure to quantify the trust of the model prediction makes use of the available information provided by the residuals of the predictions of that same model. Each of the N_p residuals (res_k^t) is stored in a pool of the most recent N_f residuals, resulting in a matrix $N_p \times N_f$. The residual at time t and for a k steps ahead prediction is defined as:

$$res_k^t = y_{k|t-k} - CGM_t \quad \forall k \in \mathbb{Z}_1^{N_p} \quad 4$$

where $y_{k|t-k}$ is the predicted glucose by the model at time t that was predicted k samples in the past. CGM_t is the current glucose values. The residuals are penalized using the forgetting function:

$$\tilde{res}_{k,i}^t = res_k^{t-i} \cdot \frac{\log(N_f - i + 1)}{\log N_f} \quad \forall i \in \mathbb{Z}_1^{N_f}, \forall k \in \mathbb{Z}_1^{N_p} \quad 5$$

where t stands for the current time, res_k^{t-i} are the progressively more distant residuals, and $\tilde{res}_{k,i}^t$ is its analog penalized residual. Note that for each timestamp t each of the residuals is penalized differently. More recent residuals are considered equally meaningful in the pool of residuals, and the relevance of older residuals decays fast when approaching the limit $N_f = 60$ samples. The normalized decay function is depicted in Figure 1. It should be noted that the depicted normalized forgetting function is merely an exemplary embodiment and both the limit and the forgetting function can be altered if required for testing of other algorithms, per the choice of the designer.

[0051] As can be seen in FIG. 1, a residual decay is depicted by plotting the decay over time, according to an exemplary embodiment.

[0052] If the calculated residuals are very large, the predictions are considered to be poor and the tuning of the controller is considered at its best when it is relaxed, i.e. conservative.

On the other hand, if the residuals are low, a more aggressive approach towards tuning the controller is preferred.

[0053] Every row (\widetilde{res}_k^t) of the penalized residuals comprises an empirical sample of the probabilistic distribution of errors of the prediction model at current time t and predicted k samples before. A good estimation of the confidence interval of the current prediction error can be calculated from the fifth and ninety-fifth percentiles, according to an exemplary embodiment, of this sample. The below formulae provide the definition of intervals from residual distribution:

$$\widehat{B}_k = \max(\text{percentile}_{95}(\widetilde{res}_k^t), 0) \quad \forall k \in \mathbb{Z}_1^{Np} \quad 6$$

$$\check{B}_k = \min(\text{percentile}_5(\widetilde{res}_k^t), 0) \quad \forall k \in \mathbb{Z}_1^{Np} \quad 7$$

where \widehat{B}_k and \check{B}_k are the empirically defined upper and lower boundaries of the confidence interval of the k steps ahead prediction error, and $\text{percentile}_x(\cdot)$ is the x -th percentile function. Both boundaries are defined strictly positive or negative in order to smooth the initialization of the algorithm. Note that both boundaries are dependent on the residual pool at the current time t , so they must be considered as time-varying values. For simplicity, the sub-index k is dropped from now on, and only the first row of the residuals pool is used for the calculation of the trust index. Similarly, only the one step-ahead prediction is considered relevant hereinafter, however the disclosure is not limited thereto.

[0054] The trust index function is designed to obtain a scalar value that comprises both the uncertainty of the previous model predictions and the coherence of the new measurement with that uncertainty. The trust index may be calculated as follows:

$$T_t := \widehat{B} - \check{B} + d_H(\text{CGM}_t, [\check{B} + y_t, \widehat{B} + y_t]) \quad 8$$

$$d_H(x, [a, b]) := \begin{cases} 0 & \text{if } x \in [a, b] \\ x - b & \text{if } x > b \\ a - x & \text{if } x < a \end{cases} \quad 9$$

where $d_H(x, [a, b])$ stands for the Hausdorff distance from a point x to an interval $[a, b]$, and y_t stands for the model's prediction of the current CGM sample. The trust index is low whenever the confidence interval is small, and it grows whenever the boundaries grow, or if the current experimental measurement falls out of the bounds. The boundaries are not fixed either, and eventually, will grow whenever enough measurements fall out of the bounds, permanently increasing the trust index (until the forgetting algorithm catches up). A

very high confidence is defined by design at $T_t \leq 5 \left[\frac{mg}{dL}\right]$ and the lowest possible confidence corresponds to a trust index value of $T_t \geq 30 \left[\frac{mg}{dL}\right]$.

[0055] FIG. 2 depicts a relationship between confidence on the current measurement and the aggressiveness of the current tuning, according to the adaptive MPC controller.

[0056] The trust index in FIG. 2 is implemented into the closed-loop scheme presented by relating it to the aggressiveness tuning of the controller, as illustrated in Figure 2.

[0057] As can be seen in FIG. 2, a high trust confidence rating corresponds to an aggressive tuning of the MPC controller and a low trust confidence rating corresponds to a conservative tuning of the MPC controller.

[0058] Every possible scenario of the correlation between the trust index and the controller aggressiveness is explained in the table depicted in FIG. 3.

[0059] FIG. 3 depicts a table describing controller behavior based on the glucose scenario and trust index value, according to an exemplary embodiment.

[0060] As can be seen in FIG. 3, under a high trust index the controller delivers insulin freely and aggressively to prevent hyperglycemia, and the controller reduces insulin delivery and may deliver basal insulin if excursions are small and predicted glucose is rising to prevent hypoglycemia. On the other hand, under a low trust index, the controller refrains on acting aggressively, but still delivers insulin to prevent hyperglycemia, and the controller shuts off the pump and may not deliver insulin even if model predictions indicate rising glucose levels to prevent hypoglycemia.

[0061] According to an exemplary embodiment, for a $T_t \leq 5$ mg/dL (high trust on the predictions) the aggressive tuning is selected, that is, $\hat{R} = 800$, $R^\vee = 500$. On the other hand, when the index surpasses the upper limit $T_t \geq 30$ mg/dL, the conservative tuning is used: $\hat{R} = 50000$, $R^\vee = 10$. Anywhere in between those thresholds both \hat{R} and R^\vee are linear functions of T_t for every time stamp t , that is, the controller gains are updated in real time as T_t is updated. This dependence of the tuning parameters on T_t , which ultimately makes them dependent on time, is what defines the aggressiveness tuning as adaptive. A graphical representation of the adaptation of both parameters, according to the exemplary embodiment, is discussed in more detail below with reference to FIG. 12.

[0062] While referencing FIGS. 4, 5 and 6 results of using a aggressively tuned, a conservatively tuned and an adaptively tuned MPC controller will be compared.

[0063] As discussed above, according to an exemplary embodiment, for a trust index value of $T_t \leq 5 \left[\frac{mg}{dL}\right]$ the aggressive tuning for the asymmetrical MPC controller is selected,

i.e. $\hat{R}=500$, $\check{R}=300$. On the other hand, when the index surpasses the upper limit $T_t \geq 30 \left[\frac{mg}{dL} \right]$, the conservative tuning is used, i.e. $\hat{R}=50000$, $\check{R}=10$. Anywhere in between both controller parameters are tuned linearly with T_t for every timestamp t .

[0064] The simulations performed using the UVA/Padova simulator aim at comparing the dynamic modulation of the tuning parameters of the MPC controller (adaptively tuned) to the two static MPC tuning states described above as “aggressive” and “conservative”.

[0065] FIG. 4 depicts a control variability grid analysis for 10 in silico subjects from the UVA/Padova metabolic simulator using an aggressive tuning of the MPC controller, according to an exemplary embodiment.

[0066] The control variability grid analysis chart for the “aggressive” controller scheme is shown in FIG. 4, with hyper- and hypo-glycemia events (E-zone) reported for 12/100 simulations, while the mean and 1-sigma ellipse falls completely within the B-zone, considered safe.

[0067] FIG. 5 depicts a control variability grid analysis for 10 in silico subjects from the UVA/Padova metabolic simulator using a conservative tuning of the MPC controller, according to an exemplary embodiment.

[0068] In FIG. 5, the control variability grid analysis chart for the “conservative” controller is shown, displaying not a single simulation in the E-zone, but 8 simulations where located in zone C and 8 in zone D. The centroid is displaced to the left with respect to the case displayed in FIG. 4, as expected from a more conservative controller.

[0069] FIG. 6 depicts a control variability grid analysis for 10 in silico subjects from the UVA/Padova metabolic simulator using an adaptive tuning of the MPC controller, according to an exemplary embodiment.

[0070] In FIG. 6, the control variability grid analysis chart for the adaptive scheme is displayed. As with the conservative case, no simulations are placed in the E-zone, but 15/100 simulations fall in the D-zone. The centroid is displaced to the left slightly less than that in the case of FIG. 5.

[0071] Metrics are reported as mean plus-minus standard deviation of the in silico population. The performance metrics for the “default,” “aggressive,” and “conservative” tuning are statistically compared to the proposed adaptive algorithm. In the bolus strategy, the adaptive controller shows very similar performance (in mean glucose and time in range) to that of the “default” controller. However, the adaptive strategy is able to reduce significantly the time spent in hypoglycemia (0.01% vs 0.0%, $p < .005$). Hypoglycemia is also rarely seen

when using the conservative controller, but at the cost of a higher average glucose (137 mg/dL vs 134 mg/dL, $p < .005$).

[0072] In the GRID scenario, the adaptive algorithm shows significantly lower average glucose than the conservative tuning (149 mg/dL vs 154 mg/dL, $p < .005$). Furthermore, the proposed algorithm shows a significant reduction in hypoglycemia with respect to the aggressive case (0.16% vs 1.04%, $p < .005$), with minimal increase in average glucose and good time in the safe glucose range. The time in hypoglycemia is not different for the conservative case and the proposed adaptive algorithm. Time in hypoglycemia for the proposed adaptive algorithm is significantly lower than for the default controller (0.16% vs 0.67%, $p < .005$), but in this case at the cost of a higher average glucose. However, the time in range for the proposed algorithm is not significantly different than that of the default tuning. Performance metrics for the above described exemplary embodiment is further discussed later, in reference to Fig. 13.

[0073] For a complete comparison, the mean glucose, time in target range, time in hypoglycemia and time in hyperglycemia for the whole cohort of the *in silico* population is depicted in the table in FIG. 7. FIG. 7 depicts a table providing a performance comparison metric of an MPC controller tuned aggressively, conservatively and adaptively, according to an exemplary embodiment. The performance metrics for both the “aggressive” and “conservative” tuning are statistically compared to the adaptive algorithm.

[0074] As can be seen in FIG. 7, the adaptive algorithm shows significantly lower average glucose than the conservative tuning. The adaptive algorithm furthermore, shows a significant reduction in hypoglycemia with respect to the aggressive case, with minimal increase in average glucose and good time in the safe glucose range. The time in hypoglycemia is not different for the conservative case and the adaptive algorithm. The time in hyperglycemia is no better for the adaptive controller than it is for the conservative case, and the aggressive controller improves over the other two cases significantly.

[0075] The control variability grid analysis chart for the aggressive tuning in Figure 4 shows a very different performance than those of the conservative and adaptive controllers (Figures 5 and 6), with multiple cases of mild and severe hypoglycemia. However, little information can be extracted by comparing the conservative and adaptive controllers. Both figures are very similar, with the centroid and the deviation ellipse being placed in a very similar position. However, when looking at FIG. 7, the adaptive tuning of the controller parameters shows better performance, clearly decreasing the average glucose from 153.7 mg/dL to 149.3 mg/dL, while increasing the time in range (70-180 mg/dL) from 71.9% to

74.7%. The numerical results also show that there is no significant difference between the time spent in hypoglycemia ($BG < 70$ mg/dL) for the conservative tuning at 0.02% and the adaptive tuning at 0.06%. Control variability grid analysis charts are usually very informative for spotting severe hypoglycemia cases and large glucose excursions. This is especially true in the postprandial periods where glucose can reach very high levels. However, the information control variability grid analysis charts provide may be insufficient to assess the performance of a long term controller.

[0076] The results for the adaptive tuning are very promising as indicated in FIG. 7, improving the performance of the existing conservative controller without increasing the risk of hypoglycemia to the patient. The time in range is still not as good as it is for the aggressive controller, but the aggressive scenario supposed a much greater danger to the patient than the adaptive case in terms of hypoglycemic risk. The exemplary embodiment of the trust index implemented in this controller allows achievement of better control for the user by good use of already available information and without any additional risk to the patient's health.

[0077] FIG. 8 depicts a chart providing a comparison in median glucose concentration and interquartile range (top) and mean insulin pump administration (bottom) between aggressive static tuning (plotted in blue - continuous line) and adaptive tuning (plotted in yellow - discontinuous line) of an MPC controller, the simulated protocol including three unannounced meals compensated using the glucose rate increase detector (GRID) meal detection algorithm, according to an exemplary embodiment.

[0078] In FIG. 8 the glucose concentration and insulin infusion of the adaptive scenario and the aggressive controller are plotted overlapping the interquartile range. The insulin bolus delivered after the meals is evident for the three unannounced meals, showing very clear spikes shortly after the meal ingestion, as the algorithm takes some time to detect each meal.

[0079] FIG. 9 depicts a chart providing a comparison in median glucose concentration and interquartile range (top) and mean insulin pump administration (bottom) between conservative static tuning (plotted in blue - continuous line) and adaptive tuning (plotted in yellow - discontinuous line) of an MPC controller, the simulated protocol including three unannounced meals compensated using the glucose rate increase detector (GRID) meal detection algorithm, according to an exemplary embodiment.

[0080] In Figure 9 the same plotting as FIG. 8 is done in a comparison between the conservative and the adaptive controllers. Again, the insulin bolus delivered after the meals is evident for the three unannounced meals, showing very clear spikes shortly after the meal ingestion, as the algorithm takes some time to detect each meal.

[0081] The median blood glucose time series shown in the top block of Figure 8 shows little to no difference in the postprandial period after the first meal of the simulation, which was the smallest of the three (50g CHO). The same conclusion can be extracted when looking at Figure 9. Glucose excursions from such a small meal may be countered by the controller independently of the tuning. On the other two meals, clear differences on the median glucose can be observed, being the aggressive glucose lower than the proposed adaptive tuning, and lower than the conservative case. This shows that in the presence of a severe hyperglycemic event, the adaptive controller is allowed to deliver more insulin than the conservative controller. It is also clear, based on the depicted exemplary embodiments, that the adaptive controller delivers this insulin at the safest time, i.e. when the model predictions are accurate, reducing the occurrence of postprandial hypoglycemic events.

[0082] Average insulin delivery, shown in the bottom plots of figures 8 and 9, displays no oscillatory behavior for the proposed algorithm or the static tuning scenarios. Average postprandial insulin delivery is lower in the adaptive tuning controller than in the aggressive static case, which explains the lower occurrence of hypoglycemia shown in FIG 7. As for the conservative versus adaptive case, there seems to be no difference on the average value of insulin delivered at any time of the simulation.

[0083] All the meals were not announced to the controller, but a meal detection algorithm based solely on the CGM measurements was implemented, according to the depicted exemplary embodiment. The action taken at the detection of each meal is displayed as large insulin spikes in the bottom plots of Figures 8 and 9. The use of a meal detection algorithm influences the performance of the controller when the meals are unannounced as it is reproduced in the simulation results presented above.

[0084] FIG. 10 depicts a block diagram of an apparatus incorporating an MPC controller, according to an exemplary embodiment.

[0085] As can be seen in FIG. 10, an exemplary embodiment of an artificial pancreas apparatus 10000 may incorporate a memory/storage 10001, and an MPC controller 10002. The MPC controller may further incorporate a processor 10003. Although a single processor is depicted in FIG. 10, the disclosure is not limited thereto and the MPC controller may incorporate a plurality of processors, according to other exemplary embodiments.

[0086] FIG. 11 depicts a flowchart describing the process to adaptively tune an MPC controller, according to an exemplary embodiment.

[0087] The process starts at step 1101 where model predictions are performed on glucose concentrations. Based on the model predictions, residuals are determined in step 1102. The

residual may be determined using formula no. 4 recited above, according to an exemplary embodiment.

[0088] Further, in step 1103, a forgetting function is used to penalize the residuals in step 1103. The penalized residual may be calculated using formula no. 5 recited above, according to an exemplary embodiment. At step 1104, upper and lower boundary of the confidence estimate in the penalized residuals are calculated. Formulae nos. 6 and 7 can be used for such a calculation according to an exemplary embodiment.

[0089] The calculated upper and lower boundary is then used to calculate the trust index in step 1105. The trust index may be calculated using formulae nos. 8 and 9 recited above, according to an exemplary embodiment.

[0090] At step 1106, the calculated trust index is used to adaptively tune the MPC controller.

[0091] A novel strategy for adaptive tuning of a zone based MPC controller has been described above and its application to a virtual cohort of people with type 1 diabetes. Repeated residuals from the model predictions are stored and a scalar trust index is deduced quantifying the uncertainty of the model predictions and the coherence of the glucose sample at each step, as further detailed using the flowchart of FIG. 11. The MPC controller is then tuned in real time based on the value of the trust index. As can be seen from the simulation comparisons depicted in FIGS. 4-9, the adaptive strategy results in better glucose control than the conservative counterpart, while producing less hypoglycemic episodes and no abnormal insulin delivery profiles were observed upon the application of the adaptive strategy. Delivery of insulin is then controlled based on the adaptively tuned MPC controller.

[0092] FIG. 12 depicts adaptation scheme for parameters \hat{R} (top) and \check{R} (bottom).

[0093] As can be seen, for trust values close to 100% (almost perfect predictions) the parameters assume the aggressive configuration ($\hat{R}=800$, $\check{R}=500$). On the other hand, at a low trust, the parameters are set to the conservative mode ($\hat{R}=50000$, $\check{R}=10$), according to an exemplary embodiment.

[0094] FIG. 13 depicts the Mean and standard deviation of the control performance metrics of the zMPC controller for the Default, Aggressive, Conservative, and Adaptive tuning, according to an exemplary embodiment. The performance metrics for the exemplary embodiments described above in the disclosure are depicted in FIG. 13.

[0095] “p” values are calculated to show significance between the proposed (adaptive) and other 3 controllers. Values with an asterisk show statistical significance ($p < .05$) between the proposed tuning and the default column, as can be seen in FIG. 13.

[0096] When comparing strategies GRID and Bolus from FIG 13, a larger disparity can be observed for the results of each controller configuration for the GRID strategy. This is a direct result of the prandial bolus dosing limiting the amount of insulin that the controller is able to utilize in the postprandial period, which is one of the safety mechanisms implemented in the controller structure. This yields, for the bolus strategy, more homogeneous results than those of the GRID strategy. As for the GRID strategy results, the adaptive controller produces significantly larger average glucose than the default controller configuration (149 mg/dL vs 147 mg/dL), according to an exemplary embodiment.

[0097] However, this difference is clinically irrelevant, with an estimated A1C (based on the average glucose) of 6.81% for the adaptive control and 6.74% for the default controller configuration, according to an exemplary embodiment. The “A1C-average glucose” regression was calculated following: $A1C = (\text{avg_glucose} + 46.7) / 28.7$. The performance difference is not significant when looking at the time in range (74.6% vs 75% $p = .12$), according to an exemplary embodiment. Furthermore, the adaptive scheme is able to significantly reduce the time spent in hypoglycemia, down to 0.16% from the 0.67% of the default controller configuration. The time spent in hypoglycemia for the adaptive scheme is not significantly different than that achieved by the conservative controller configuration. Similarly to what was observed for the bolus strategy experiment, the conservative controller produces a higher average glucose than that of the adaptive control, while also showing a significantly lower time in range. The aggressive controller is able to lower the average glucose more than any of the other configurations, at the cost of much higher frequency of hypoglycemia (1%), according to an exemplary embodiment.

[0098] The present disclosure is not limited to the precise construction and compositions disclosed herein; any and all modifications, changes, and variations apparent from the foregoing descriptions are within the spirit and scope of the disclosure as defined in the appended claims. Moreover, the present concepts expressly include any and all combinations and sub combinations of the preceding elements and aspects. An implementation of an apparatus that falls within the inventive concept does not necessarily achieve any of the possible benefits outlined above: such benefits are dependent on the specific use case and specific implementation, and the possible benefits mentioned above are simply examples.

[0099] Although the concepts have been described above with respect to the various embodiments, it is noted that there can be a variety of permutations and modifications of the described features by those who are familiar with this field, only some of which have been presented above, without departing from the technical ideas and scope of the features, which is defined by the appended claims.

[00100] Further, while this specification contains many features, the features should not be construed as limitations on the scope of the disclosure or the appended claims. Certain features described in the context of separate embodiments can also be implemented in combination. Conversely, various features described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable sub-combination.

[00101] Although the drawings describe operations in a specific order and/or show specific arrangements of components, one should not interpret that such specific order and/or arrangements are limited, or that all the operations performed and the components disclosed are needed to obtain a desired result. There are numerous hardware and software devices that can be configured to forward data units in the manner described in the present disclosure with respect to various embodiments. Accordingly, other implementations are within the scope of the following claims.

CLAIMS

1. A method of adaptively tuning a Model Predictive Control (MPC) controller, using at least one processor, the method comprising:
 - determining, using at least one of said at least one processor, a plurality of residuals, wherein each of the plurality of residuals is based on a difference between a measured glucose value and a glucose value predicted by a model at a previous time among of a plurality of time samples in the past;
 - storing, in a memory, the determined residuals;
 - calculating a trust index, using at least one of said at least one processor, by quantifying uncertainty of the model using the stored residuals;
 - tuning the MPC controller, in real time, using at least one of said at least one processor, based on the calculated value of the trust index; and
 - determining an amount of insulin to be delivered by an insulin pump based on insulin amounts determined by the tuned MPC controller.

2. The method of claim 1, wherein
 - the determining further comprises determining the residual at time t and for a prediction k steps ahead using the formula:

$$res_k^t = y_{k|t-k} - CGM_t \quad \forall k \in \mathbb{Z}_1^{N_p}$$

where $y_{k|t-k}$ represents predicted glucose by the model at time t that was predicted k samples in the past, N_p represents prediction horizon and CGM_t represents current glucose values.

3. The method of claim 1 or 2, wherein the storing further comprises storing each of the N_p residuals (res_k^t) in a pool of most recent N_f residuals, resulting in a matrix $N_p \times N_f$.

4. The method according to any one of claims 1 to 3, further comprising:
 - penalizing, using at least one of said at least one processor, the determined residuals using a forgetting function, wherein

the forgetting function is determined using the formula:

$$\widetilde{res}_{k,i}^t = res_k^{t-i} \cdot \frac{\log(N_f - i + 1)}{\log N_f} \quad \forall i \in \mathbb{Z}_1^{N_f}, \forall k \in \mathbb{Z}_1^{N_p}$$

where t represents current time, res_k^{t-i} represent progressively more distant residuals, and $\widetilde{res}_{k,i}^t$ represents the penalized residual.

5. The method of claim 4, wherein for each timestamp t each residual is penalized differently.
6. The method according to any one of claims 1 to 5, wherein tuning the controller comprises decreasing an aggressiveness of the MPC controller if the trust index is below a threshold.
7. The method according to any one of claims 1 to 5, wherein tuning the controller comprises increasing an aggressiveness of the MPC controller if the trust index is above a threshold.
8. The method of claim 4 or 5, further comprising:

calculating, using at least one of said at least one processor, an estimate of a confidence interval of a current prediction error based on the penalized residuals using the formula:

$$B_x = \max(\text{percentile}_x(\widetilde{res}_k^t), 0) \quad \forall k \in \mathbb{Z}_1^{N_p}$$

where B_k is an empirically defined boundary of the confidence interval of a prediction error k steps ahead, and $\text{percentile}_x(\widetilde{res}_k^t)$ represents x -th percentile function of the penalized residual at time t and for a prediction k steps ahead.

9. The method of claim 8, wherein the calculating the estimate of the confidence interval comprises calculating an empirically defined upper boundary \widehat{B}_k and an empirically defined lower boundary \check{B}_k using the formula:

$$\begin{aligned} \widehat{B}_k &= \max(\text{percentile}_{95}(\widetilde{res}_k^t), 0) & \forall k \in \mathbb{Z}_1^{N_p} \\ \check{B}_k &= \min(\text{percentile}_5(\widetilde{res}_k^t), 0) & \forall k \in \mathbb{Z}_1^{N_p} \end{aligned}$$

10. The method of claim 9, wherein the calculating the trust index comprises calculating the trust index based on the calculated empirically defined upper boundary \widehat{B}_k and the calculated empirically defined lower boundary \check{B}_k using the formula:

$$T_t := \widehat{B}_k - \check{B}_k + d_H(CGM_t, [\check{B} + y_t, \widehat{B} + y_t])$$

$$d_H(x, [a, b]) := \begin{cases} 0 & \text{if } x \in [a, b] \\ x - b & \text{if } x > b \\ a - x & \text{if } x < a \end{cases}$$

where $d_H(x, [a, b])$ represents Hausdorff distance from a point x to an interval $[a, b]$, and y_t represents model's prediction of current CGM sample.

11. A method of insulin delivery using an artificial pancreas, the artificial pancreas including at least one processor, the method comprising:
 - determining, using at least one of said at least one processor, a plurality of residuals, wherein each of the plurality of residuals is based on a difference between a measured glucose value and a glucose value predicted by a model at a previous time among of a plurality of time samples in the past;
 - storing, in a memory, the determined residuals;
 - calculating a trust index, using at least one of said at least one processor, by quantifying uncertainty of the model using the stored residuals;
 - tuning a model prediction control (MPC) controller, in real time, using at least one of said at least one processor, based on the calculated value of the trust index; and
 - controlling, using at least one of said at least one processor, the delivery of insulin based on insulin amounts determined by the tuned MPC controller.

12. An artificial pancreas apparatus for insulin delivery, the artificial pancreas comprising:
 - at least one non-transitory memory operable to store program code;
 - a model prediction control (MPC) controller including at least one processor operable to read said program code and operate as instructed by said program code, said program code causing the at least one processor to:
 - determine a plurality of residuals, wherein each of the plurality of residuals is based on a difference between a measured glucose value and a glucose value predicted by a model at a previous time among of a plurality of time samples in the past;
 - store, in the at least one non-transitory memory, the determined residuals;
 - calculate a trust index by quantifying uncertainty of the prediction models using the stored residuals;
 - tune the MPC controller in real time based on the calculated value of the trust index;
 - and

determine an amount of insulin to be delivered by an insulin pump of the artificial pancreas based on the tuned MPC controller.

13. The artificial pancreas apparatus of claim 12, wherein tuning the MPC controller, in real time, based on the calculated value of the trust index comprises increasing the aggressiveness of the MPC controller if the trust index is above a threshold and decreasing the aggressiveness of the MPC controller if the trust index is below the threshold.

FIG. 1




FIG. 2




FIG. 3

Controller behavior for every glucose scenario and each trust index value

	High Trust Index	Low Trust Index
Hyperglycemia	Controller delivers insulin freely and aggressively	Controller refrains on acting aggressively, but still delivers insulin
Hypoglycemia	Controller reduces insulin delivery. Controller may deliver basal insulin if excursions are small and predicted glucose is rising	Controller shuts off the pump. May not deliver insulin even if model predictions indicate rising glucose levels

FIG. 4




FIG. 5




FIG. 6




FIG. 7

Performance metrics of the controller for aggressive, conservative and dynamic tuning. p-values shown for comparison against the adaptive tuning

Tuning	Aggressive	Conservative	Adaptive
Average glucose [mg/dL (p-value)]	142.9 (<0.005)	153.7 (<0.005)	149.3
Time in range 70-180 [% (p-value)]	77.3 (<0.005)	71.9 (<0.005)	74.7
Time in hypo < 70 [% (p-value)]	0.98 (<0.005)	0.02 (0.38)	0.06
Time in hyper > 300 [% (p-value)]	0.69 (<0.005)	1.48 (0.34)	1.4

FIG. 8




FIG. 9




FIG. 10


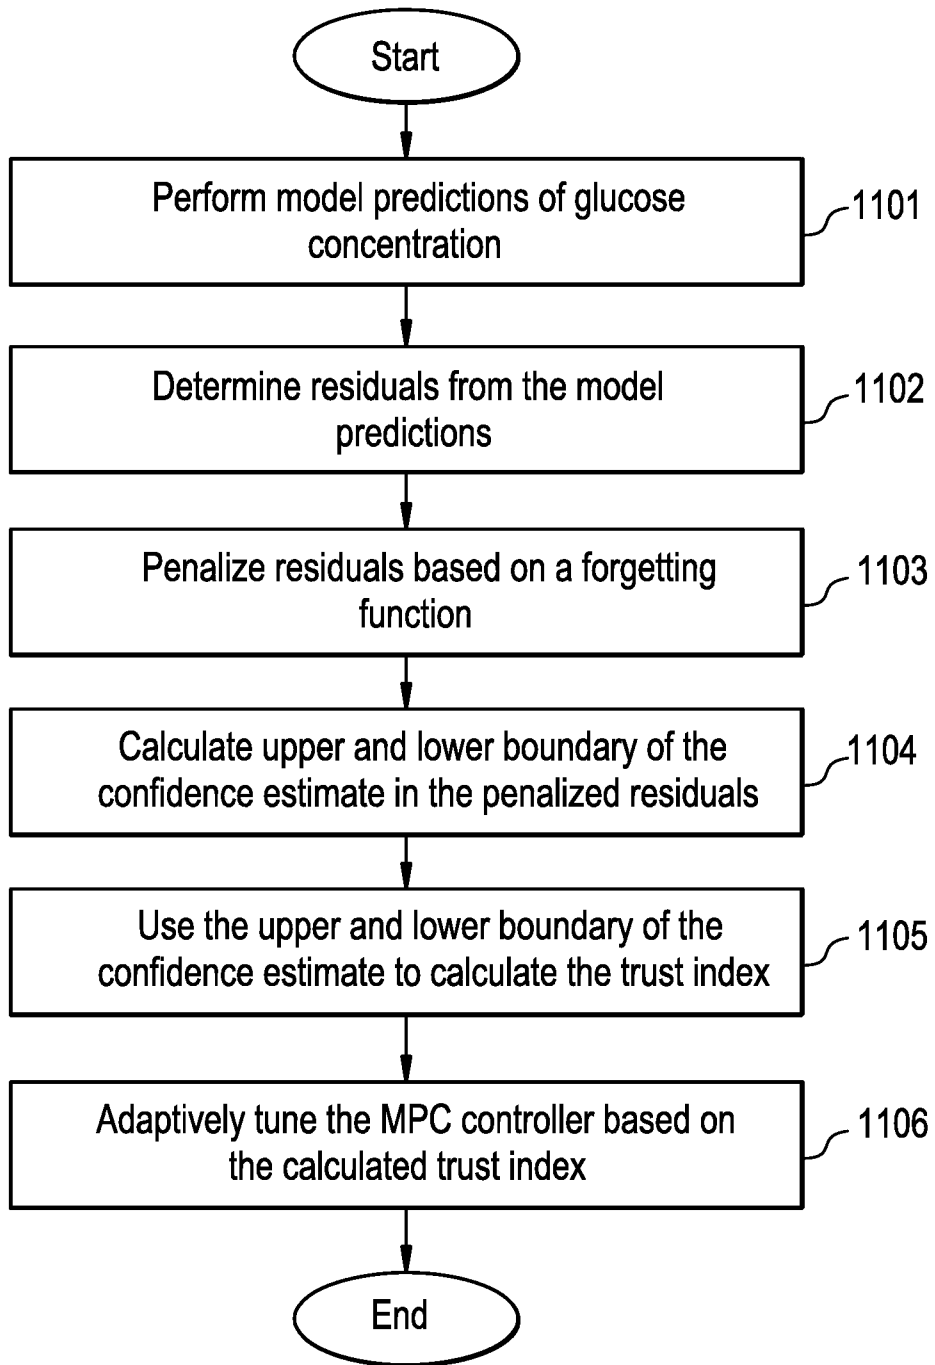


FIG. 11



11/12




FIG. 12

12/12

FIG. 13

Tuning	Adaptive	Default	Aggressive	Conservative
Bolus strategy				
Average glucose [mg/dL (p value)]	134 ± 6	134 ± 6 (.2)	131 ± 6 (<.005)	137 ± 7 (<.005)
Total insulin [IU (p value)]	24.9 ± 5.6	24.8 ± 5.5 (.33)	25.5 ± 5.8 (.0057)	24.4 ± 5.5 (.03)
Time in target range 70-180 [% (p value)]	91 ± 5.6	90.7 ± 5.9 (.16)	92.3 ± 6.1 (<.005)	89.8 ± 5.8 (<.005)
Time in hypoglycemia <70 [% (p value)]	0 ± 0*	0.01 ± 0.16 (<.005)	0.1 ± 0.57 (<.005)	0 ± 0 (1)
Time in hyperglycemia >300 [% (p value)]	0.01 ± 0.13	0.01 ± 0.22 (.25)	0.01 ± 0.21 (.22)	0 ± 0 (.08)
GRID strategy				
Average glucose [mg/dL (p value)]	149 ± 10*	147 ± 8 (<.005)	143 ± 7 (<.005)	154 ± 13 (<.005)
Total insulin [IU (p value)]	23.4 ± 5.2*	23.8 ± 5.1 (.03)	24.5 ± 5.8 (<.005)	22.8 ± 5.1 (0.1)
Time in target range 70-180 [% (p value)]	74.6 ± 6.9	75 ± 7.1 (.12)	77 ± 7.2 (<.005)	71.2 ± 9.4 (<.005)
Time in hypoglycemia <70 [% (p value)]	0.16 ± 0.84*	0.67 ± 1.82 (<.005)	1.04 ± 2.43 (<.005)	0.15 ± 1 (.44)
Time in hyperglycemia >300 [% (p value)]	1.07 ± 2.84	1.02 ± 2.31 (.34)	0.7 ± 1.62 (<.005)	1.37 ± 3.5 (.02)