



(51) International Patent Classification:

Not classified

(21) International Application Number:

PCT/US2024/034051

(22) International Filing Date:

14 June 2024 (14.06.2024)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

63/508,235

14 June 2023 (14.06.2023)

US

(71) Applicant: ARIZONA BOARD OF REGENTS ON BEHALF OF ARIZONA STATE UNIVERSITY [US/US]; 699 S. Mill Avenue, Suite 601, Room 691AA, Tempe, Arizona 85281 (US).

(72) Inventors: ZADEH, Hassan; c/o Arizona Board of Regents on Behalf of Arizona State University, 699 S. Mill Avenue, Suite 601, Room 691AA, Tempe, Arizona 85281 (US). AREFEEN, Asiful; c/o Arizona Board of Regents on Behalf of Arizona State University, 699 S. Mill Avenue,

Suite 601, Room 691AA, Tempe, Arizona 85281 (US). KHAMESIAN, Saman; c/o Arizona Board of Regents on Behalf of Arizona State University, 699 S. Mill Avenue, Suite 601, Room 691AA, Tempe, Arizona 85281 (US).

(74) Agent: BAI, Ari et al.; POLSINELLI PC, PO BOX 140310, Kansas City, Missouri 64114-0310 (US).

(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, CV, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IQ, IR, IS, IT, JM, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, MG, MK, MN, MU, 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, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

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

(54) Title: SYSTEMS AND METHODS FOR MODELING AND SIMULATING GLYCEMIC RESPONSE FOR BEHAVIORAL LIFESTYLE INTERVENTIONS




FIG. 1

(57) Abstract: A system implementing a deep learning-based blood glucose simulator for glycemic response takes data from multi-modal information sources and allows users to virtually examine the effects of changes in behavioral parameters on blood glucose. The system predicts blood glucose values based on observation data and iteratively simulates intervention strategies to identify an action for blood glucose management.

SZ, 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, ME, 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).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

SYSTEMS AND METHODS FOR MODELING AND SIMULATING GLYCEMIC RESPONSE FOR BEHAVIORAL LIFESTYLE INTERVENTIONS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This is a U.S. Non-Provisional Patent Application that claims benefit to U.S. Provisional Patent Application Serial No. 63/508,235 filed 14 June 2023, which is herein incorporated by reference in its entirety.

GOVERNMENT SUPPORT

[0002] This invention was made with government support under 2210133 awarded by the National Science Foundation. The government has certain rights in the invention.

FIELD

[0003] The present disclosure generally relates to synthetic data generation, and in particular, to a system and associated method for modeling and simulating glycemic responses to intervention strategies.

BACKGROUND

[0004] Glucose control is critical in both diabetes prevention and diabetes management. Poor glucose control leads to abnormal events such as hyperglycemia and hypoglycemia. Regular exposure to dysglycemia can increase the risk of complications, including cardiovascular disease, reduced eyesight, cancer, impaired glycemic control, and seizure. CGM sensors are commonly used to monitor blood glucose. However, their utility for behavior change is currently limited. Current CGM systems lack computational abilities to actively influence users' behavior and work as a decision support for individuals with or at risk for diabetes. Such a limitation thwarts the effectiveness of CGM systems to proactively promote health enhancing behaviors associated with food consumption, medication, sleep, and physical activity, to maintain stable blood glucose levels.

[0005] It is with these observations in mind, among others, that various aspects of the present disclosure were conceived and developed.

SUMMARY

[0006] A system for forecasting a blood glucose value and informing a user of an intervention action to avoid dysglycemia based on the blood glucose value includes a processor in communication with a memory and one or more sensor devices, the memory including instructions executable by the processor to: access, by a processor in communication with a memory and one or more sensor devices, a set of multimodal observation data for a first time interval, the set of multimodal observation data including values of a plurality of parameters that affect values of a factor of interest and the set of multimodal observation data including time-dependent values for the factor of interest; generate, by application of the set of multimodal observation data as input to a prediction model, a sequence of prediction values for the factor of interest across a prediction horizon interval, the prediction horizon interval including a second time interval following the first time interval associated with the set of multimodal observation data; and determine, over one or more iterations, an intervention action to be applied at an intervention time step within the second time interval with respect to the plurality of parameters that is predicted to result in an intervention sequence of prediction values for the factor of interest having a predicted maximum value and a predicted minimum value that are within a target range.

[0007] In some examples, the factor of interest includes a blood glucose level or CGM value, and the time-dependent values of the factor of interest for the first time interval are measurable by a blood glucose sensor of the one or more sensor devices. The plurality of parameters can include a consumed carbohydrate amount. The plurality of parameters can also include one or more of a basal insulin dosage and a bolus insulin dosage, and the time-dependent values of the basal insulin dosage or the bolus insulin dosage for the first time interval being reported by an insulin pump device in communication with the processor. Further, the set of multimodal observation data associated with the first time interval can include one or more physiological signals. The physiological signals can include, but are not limited to: a skin temperature measurable by a temperature sensor of the one or more sensor devices; a heart rate measurable by a heart rate sensor of the one or more sensor devices; and a galvanic skin response value measurable by a galvanic skin response sensor of the one or more sensor devices. Further, the set of

multimodal observation data associated with the first time interval can include events such as sleep or physical activity which may be detected by a user device or reported by a user.

[0008] The memory can further include instructions executable by the processor to: evaluate, based on the sequence of prediction values for the factor of interest associated with the set of multimodal observation data, whether a predicted maximum value or a predicted minimum value of the sequence of prediction values is outside of the target range.

[0009] Further, the system can include a display device in communication with the processor. The memory can further include instructions executable by the processor to: display a graphical representation at the display device that illustrates the set of multimodal observation data for the first time interval and the sequence of prediction values for the factor of interest across the prediction horizon interval.

[0010] The memory can further include instructions executable by the processor to: generate, for an iteration of the one or more iterations, an intervention set of multimodal observation data that includes a modified value of a parameter of the plurality of parameters for the intervention time step within the second time interval, the modified value corresponding to the intervention action; generate, for the iteration of the one or more iterations and by application of the intervention set of multimodal observation data as input to the prediction model, the intervention sequence of prediction values for the factor of interest following the intervention time step; and evaluate, for the iteration of the one or more iterations and based on the sequence of prediction values for the factor of interest within the prediction horizon interval, whether the predicted maximum value or the predicted minimum value of the intervention sequence of prediction values is outside of the target range.

[0011] Further, the memory can further include instructions executable by the processor to: display a graphical representation at the display device that illustrates the intervention set of multimodal observation data and the intervention sequence of prediction values for the factor of interest.

[0012] For implementation of the prediction model, the memory can further include instructions executable by the processor to: apply a sliding window to the set of multimodal observation data resulting in an input feature sequence; extract, by a convolutional layer of the prediction model, a high-level feature

sequence based on the input feature sequence; and generate, by a long short term memory layer of the prediction model, the sequence of prediction values for the factor of interest across the prediction horizon interval based on the high-level feature sequence.

[0013] For training the prediction model, the memory can further include instructions executable by the processor to: optimize a loss function that more heavily penalizes errors associated with a sequence of prediction values when the ground truth sequence of values associated with the sequence of the prediction values are within a specified region. For example, a CGM range may include a hyperglycemic range, a hypoglycemic range, and a normal range, where it is more dangerous to be “wrong” when the ground truth sequence of values (the actual blood glucose values) are in the hyperglycemic range or the hypoglycemic range. As such, the loss function may assign penalties for errors that are associated with the hyperglycemic range or the hypoglycemic range. In some examples, the assigned penalty can double for errors within fiducial regions.

[0014] For simulating potential interventions using the prediction model, the memory can further include instructions executable by the processor to: modify the input feature sequence based on an intervention set of multimodal observation data; extract, by the convolutional layer of the prediction model, an intervention high-level feature sequence based on the input feature sequence; and generate, by the long short term memory layer of the prediction model, the intervention sequence of prediction values for the factor of interest based on the intervention high-level feature sequence.

[0015] In some examples, the memory can further include instructions executable by the processor to: identify an optimal intervention action that is expected to result in an intervention sequence of prediction values that are within the target range, where more than one intervention action is predicted to result in the intervention sequence of prediction values having the predicted maximum value and the predicted minimum value within the target range.

[0016] In a further aspect, a method of forecasting a blood glucose value and informing a user of an intervention action to avoid dysglycemia based on the blood glucose value includes: accessing, by a processor in communication with a memory and one or more sensor devices, a set of multimodal observation data for a first time interval, the set of multimodal observation data including values of a

plurality of parameters that affect values of a factor of interest and the set of multimodal observation data including time-dependent values for the factor of interest; generating, by application of the set of multimodal observation data as input to a prediction model, a sequence of prediction values for the factor of interest across a prediction horizon interval, the prediction horizon interval including a second time interval following the first time interval associated with the set of multimodal observation data; and determining, over one or more iterations, an intervention action to be applied at an intervention time step within the second time interval with respect to the plurality of parameters that is predicted to result in an intervention sequence of prediction values for the factor of interest having a predicted maximum value and a predicted minimum value that are within a target range.

[0017] The plurality of parameters can include a consumed carbohydrate amount. The plurality of parameters can also include one or more of a basal insulin dosage and a bolus insulin dosage, and the time-dependent values of the basal insulin dosage or the bolus insulin dosage for the first time interval being reported by an insulin pump device in communication with the processor. Further, the set of multimodal observation data associated with the first time interval can include one or more physiological signals. The physiological signals can include, but are not limited to: a skin temperature measurable by a temperature sensor of the one or more sensor devices; a heart rate measurable by a heart rate sensor of the one or more sensor devices; and a galvanic skin response value measurable by a galvanic skin response sensor of the one or more sensor devices. Further, the set of multimodal observation data associated with the first time interval can include events such as sleep or physical activity which may be detected by a user device or reported by a user.

[0018] The method can further include: evaluating, based on the sequence of prediction values for the factor of interest associated with the set of multimodal observation data, whether a predicted maximum value or a predicted minimum value of the sequence of prediction values is outside of the target range.

[0019] The method can further include: displaying a graphical representation at a display device in communication with the processor that illustrates the set of multimodal observation data for the first time interval and the

sequence of prediction values for the factor of interest across the prediction horizon interval.

[0020] The method can further include: generating, for an iteration of the one or more iterations, an intervention set of multimodal observation data that includes a modified value of a parameter of the plurality of parameters for the intervention time step within the second time interval, the modified value corresponding to the intervention action; generating, for the iteration of the one or more iterations and by application of the intervention set of multimodal observation data as input to the prediction model, the intervention sequence of prediction values for the factor of interest following the intervention time step; and evaluating, for the iteration of the one or more iterations and based on the sequence of prediction values for the factor of interest within the prediction horizon interval, whether the predicted maximum value or the predicted minimum value of the intervention sequence of prediction values is outside of the target range.

[0021] The method can further include displaying a graphical representation at a display device in communication with the processor that illustrates the intervention set of multimodal observation data and the intervention sequence of prediction values for the factor of interest.

[0022] For implementation of the prediction model, the method can further include: applying a sliding window to the set of multimodal observation data resulting in an input feature sequence; extracting, by a convolutional layer of the prediction model, a high-level feature sequence based on the input feature sequence; and generating, by a long short term memory layer of the prediction model, the sequence of prediction values for the factor of interest across the prediction horizon interval based on the high-level feature sequence.

[0023] For training the prediction model, the method can further include: optimizing a loss function that more heavily penalizes errors associated with a sequence of prediction values when the ground truth sequence of values associated with the sequence of the prediction values are within a specified region. For example, a CGM range may include a hyperglycemic range, a hypoglycemic range, and a normal range, where it is more dangerous to be “wrong” when the ground truth sequence of values (the actual blood glucose values) are in the hyperglycemic range or the hypoglycemic range. As such, the loss function may assign penalties for errors that are associated with the hyperglycemic range or the

hypoglycemic range. In some examples, the assigned penalty can double for errors within fiducial regions.

[0024] For simulating potential intervention actions, the method can further include: modifying the input feature sequence based on an intervention set of multimodal observation data; extracting, by the convolutional layer of the prediction model, an intervention high-level feature sequence based on the input feature sequence; and generating, by the long short term memory layer of the prediction model, the intervention sequence of prediction values for the factor of interest based on the intervention high-level feature sequence.

[0025] In some examples, the method can further include: identifying an optimal intervention action that is expected to result in an intervention sequence of prediction values that are within the target range, where more than one intervention action is predicted to result in the intervention sequence of prediction values having the predicted maximum value and the predicted minimum value within the target range.

[0026] Further, a non-transitory computer readable medium can include instructions which are executable by a processor to: access, by a processor in communication with a memory and one or more sensor devices, a set of multimodal observation data for a first time interval, the set of multimodal observation data including values of a plurality of parameters that affect values of a factor of interest and the set of multimodal observation data including time-dependent values for the factor of interest; generate, by application of the set of multimodal observation data as input to a prediction model, a sequence of prediction values for the factor of interest across a prediction horizon interval, the prediction horizon interval including a second time interval following the first time interval associated with the set of multimodal observation data; and determine, over one or more iterations, an intervention action to be applied at an intervention time step within the second time interval with respect to the plurality of parameters that is predicted to result in an intervention sequence of prediction values for the factor of interest having a predicted maximum value and a predicted minimum value that are within a target range.

[0027] In some examples, the factor of interest includes a blood glucose level or CGM value, and the time-dependent values of the factor of interest for the first time interval are measurable by a blood glucose sensor of the one or

more sensor devices. The plurality of parameters can include a consumed carbohydrate amount. The plurality of parameters can also include one or more of a basal insulin dosage and a bolus insulin dosage, and the time-dependent values of the basal insulin dosage or the bolus insulin dosage for the first time interval being reported by an insulin pump device in communication with the processor. Further, the set of multimodal observation data associated with the first time interval can include one or more physiological signals. The physiological signals can include, but are not limited to: a skin temperature measurable by a temperature sensor of the one or more sensor devices; a heart rate measurable by a heart rate sensor of the one or more sensor devices; and a galvanic skin response value measurable by a galvanic skin response sensor of the one or more sensor devices. Further, the set of multimodal observation data associated with the first time interval can include events such as sleep or physical activity which may be detected by a user device or reported by a user.

BRIEF DESCRIPTION OF THE DRAWINGS

[0028] The present patent or application file contains at least one drawing executed in color. Copies of this patent or patent application publication with color drawing(s) will be provided by the Office upon request and payment of the necessary fee.

[0029] FIG. 1 is a simplified block diagram showing a system (“GlySim”) for predicting a value based on multimodal observation data and determining an intervention action to improve the value through an iterative simulation process;

[0030] FIG. 2 is a simplified diagram showing an example computing system for implementation of the system of FIG. 1;

[0031] FIG. 3 is a diagram showing an overview of a predictive model of the system of FIG. 1, which integrates data from smartphone, smartwatch, wearable CGM and insulin pump, and devises a multi-task learning model to predict glucose response;

[0032] FIG. 4 is a diagram showing a stacked CNN-LSTM architecture of the predictive model of FIGS. 1 and 3 where inputs to the model include 3-hour long multimodal data;

[0033] FIGS. 5A and 5B are diagrams showing alternative architecture layouts for the predictive model of FIGS. 1 and 3;

[0034] FIG. 6 shows an example graphical user interface for simulating glucose response based on behavioral parameters;

[0035] FIGS. 7A-7D are a series of graphical representation showing an example where the system of FIG. 1 predicts hyperglycemia, detects that increasing prior BGL is forcing predicted BGL up, and reduces prior BGL iteratively until simulation shows normoglycemia;

[0036] FIGS. 8A and 8B are a pair of graphical representations showing CGM prediction and actual for two different embodiments of the predictive model of FIGS. 1 and 3, where the predictive model of FIG. 8B incorporates a custom loss function that penalizes errors in hyper or hypoglycemic estimations; and

[0037] FIGS. 9A-9C are a series of graphical representations showing actual vs. predicted CGM values for three patients, where prediction was generated using the predictive model of FIGS. 1 and 3.

[0038] Corresponding reference characters indicate corresponding elements among the view of the drawings. The headings used in the figures do not limit the scope of the claims.

DETAILED DESCRIPTION

[0039] Effective prevention and management of diabetes relies on maintaining a normal blood glucose level, thus avoiding abnormal events such as hyperglycemia and hypoglycemia. Predicting anomalous events beforehand can potentially help patients and caregivers intervene to prevent such events through modifiable behaviors such as exercise, diet, and medication. Although Continuous Glucose Monitor (CGM) sensors have been used to monitor and forecast blood glucose level, current research lacks a computational approach that recommends a behavioral intervention to bring the glucose level to a normal range. To address this shortcoming, the present disclosure outlines GlySim, a CGM simulator that uses multimodal data to not only forecast future glucose readings but also enable a user to examine the impacts of behavior change on glucose response in advance. GlySim creates opportunities for change in food consumption, medication, and physical activity to avoid dysglycemia by pinpointing factors that cause anomalous events using Grad-CAM (Gradient-weighted Class Activation Mapping) and allowing users to observe how adjusting a behavioral factor changes glucose trajectories. GlySim is validated on a dataset of 10 patients with type 1 diabetes and achieve an overall

mean absolute error (MAE) as low as 16.5 mg/dl in simulating glycemic response. Furthermore, Glysime detects hyperglycemic events with 0.89 average precision.

[0040] In recent years, researchers have introduced machine learning algorithms to predict CGM values, but these algorithms often fall short in several areas:

- a. Their error rates (RMSE or MAE) are not sufficiently low.
- b. They fail to accurately predict hyperglycemic or hypoglycemic events when data fluctuates.
- c. They cannot accurately forecast real values for the next hour.
- d. They have not been fine-tuned to achieve optimal results.

[0041] Recent work developed algorithms to predict blood glucose from CGM sensors. One study used a moving window on blood glucose and accelerometer signals to predict future blood glucose levels with a LSTM (long short-term memory) network. Another study developed a lightweight LSTM model, aiming to deploy it on a microcontroller for enhanced practicality. Additionally, a low-cost and low-power wearable system was designed with attention-based evidential recurrent neural network convenient for use in clinical settings. Another study used multimodal data in a stacked LSTM model followed by Kalman filter to predict blood glucose values. Furthermore, a modified ResNet architecture was introduced to predict blood glucose levels 30 minutes into the future. Although the value of accurately forecasting blood glucose level can be perceived from these studies, little effort is given in designing a mechanism that enables interventions to prevent abnormal blood glucose levels such as hyperglycemia and hypoglycemia.

[0042] The present disclosure addresses these shortcomings with *GlySim*, a simulator for glycemic response that takes data from multi-modal information sources and allows users to virtually examine the effects of changes in behavioral parameters on blood glucose. Developing methods for modeling and simulating glycemic response is central to the development of a digital twin for human health. Beyond effective diabetes prevention and management, a CGM simulator will provide a platform for developing and testing novel algorithms and techniques for glucose management without extensive clinical studies. Such a simulator will also be instrumental for developing a reinforcement learning aided insulin optimizer, where the reward function is based on maintaining a normal blood

glucose level. Simulating glucose response, researchers and developers can test the effectiveness of different interventions, algorithms, and decision-making strategies in a controlled and flexible virtual environment. Moreover, GlySim enables the construction of large-scale and diverse datasets to educate and train healthcare professionals, patients, caregivers and AI models on the interpretation of CGM data in a risk-free environment.

[0043] GlySim can be viewed as a deep learning-based blood glucose simulator that enables predictive conducting of virtual interventions. In designing GlySim, a neural network is first introduced to predict future glucose readings over a prediction horizon and to identify occurrences of impending dysglycemia. A dashboard simulator is then designed to visualize the predicted signals and allow the user to tune behavioral parameters and observe the impact on the glucose curve. An intervention strategy is finally recommended by using gradients of each target prediction to determine what behavioral factor is most responsible for the predicted outcome, followed by an iterative process to virtually determine the intervention impact.

[0044] GlySim includes modules for prediction and intervention, which respond to the modeling and simulation requirement of a digital twin, respectively.

[0045] For some continuous signals, there are regions where it is more important to minimize errors. One approach is to create features and modify a predictive model to treat these sensitive regions differently. For example, in type 1 diabetes, when the CGM value is in hyperglycemia or hypoglycemia regions, patients need to be careful and take action to bring their glucose levels back to the normal range. As such, GlySim further includes several features to address these, including: classifying CGM values based on hyperglycemia and hypoglycemia thresholds to distinguish these regions; designing a custom loss function that emphasizes the significance of these classes by applying additional penalties; and fine-tuning hyperparameters.

System Overview

[0046] Referring to FIG. 1, a system **100** forecasts a factor of interest (such as a blood glucose value) and informs a user of an intervention action based on the forecasted values. The system **100** can be implemented at a computing device **200** (shown in FIG. 2 and discussed in a later section herein, which can

include a processor **220** in communication with a memory **240**) in communication with one or more sensor devices **102**.

[0047] The system **100** constructs or otherwise accesses a set of multimodal observation data **12**, discussed in further detail herein with respect to FIG. **3**. In some examples, the system **100** includes a data fusion element **110** that can perform various pre-processing operations to construct the set of multimodal observation data **12**. A subset of the set of multimodal observation data **12** can be obtained through sensor devices **102**.

[0048] To predict values of the factor of interest, the system **100** includes a prediction model **120** which accesses the set of multimodal observation data **12** (see circle (1) of FIG. **1**), and generates a sequence of prediction values **14** which correspond to forecasted values of the factor of interest. Based on the sequence of prediction values **14**, the system **100** can identify an intervention action that, if taken by a user, may improve values of the factor of interest, FIG. **1** shows an “intervention” block **130** which represents an iterative simulation process that the system **100** can apply to determine the intervention action.

[0049] At circle (2) of FIG. **1**, the system **100** can evaluate the sequence of prediction values **14** to see if an intervention action is warranted. If the sequence of prediction values **14** are sub-optimal, such as if they are outside of a target range, then at circle (3) the system **100** modifies one or more parameter values to generate an *intervention* set of multimodal observation data **22**. The intervention set of multimodal observation data **22** can correspond to a hypothetical change in behavior or choice that may be made by the user to avoid impending problems associated with sub-optimal values of the factor of interest.

[0050] At circle (4), the system **100** applies the intervention set of multimodal observation data **22** as input to the prediction model **120** to simulate what would happen if the user were to take an intervention action corresponding to the intervention set of multimodal observation data **22**. The prediction model **120** produces an *intervention* sequence of prediction values **24** that corresponds with the intervention set of multimodal observation data **22**. At circle (5), the system **100** can re-evaluate the intervention sequence of prediction values **24** to see if further modification of parameter values are necessary. Steps (3)-(5) can be iteratively repeated until a satisfactory or optimal intervention action is found.

[0051] The system **100** can further include a graphical user interface **140** (“GUI”) for display at a display device **230** (FIG. **2**) that can display graphical representations associated with the set of multimodal observation data **12**, the sequence of prediction values **14**, the intervention set of multimodal observation data **22** associated with an intervention action, and/or the intervention sequence of prediction values **24** associated with an intervention action. In examples corresponding to blood glucose monitoring, a user can view and interact with the GUI **140** to aid in understanding of how to best manage their blood glucose and how their actions and choices can be expected to affect their blood glucose.

Observation Data

[0052] With additional reference to FIG. **3**, the system **100** (in particular, the processor **220** of the computing device **200**) can access the set of multimodal observation data **12** which can be tracked over a first time interval (e.g., the last three hours or so). The set of multimodal observation data **12** can include physiological signal data (e.g., heart rate, skin temperature, galvanic skin response information, etc.) in addition to values of a plurality of parameters including carbohydrate intake, basal insulin dosage and bolus insulin dosage. Further, the set of multimodal observation data **12** can include events including sleep information, exercise (physical activity) information, and work information.

[0053] The set of multimodal observation data **12** can also include time-dependent values associated with a *factor of interest* (such as, but not limited to blood glucose level as measured using continuous glucose monitoring data (“CGM” data)), which can be affected by values of other types of information present within the set of multimodal observation data.

[0054] Some instances of the set of multimodal observation data **12** can be obtained through the one or more sensor devices **102**, e.g., the blood glucose level can be measured by a blood glucose sensor, the basal and/or bolus insulin dosages can be reported by an insulin pump device, and the physiological signal data can be measured by other sensor devices that may be commercially available, etc.

[0055] In other examples where the factor of interest is not directly observable, the system **100** may employ intermediate methods for inferring values of the factor of interest based on the set of multimodal observation data.

[0056] In some examples, the system **100** can apply one or more data fusion operations to the set of multimodal observation data **12**, e.g., at data fusion element **110**. The one or more data fusion operations can include pre-processing, temporal alignment, and/or feature extraction elements.

A. Prediction Modeling Using Deep Learning

[0057] Let $X = \{x^{(1)}, x^{(2)}, \dots, x^{(d)}\}$ be the set of d feature/sensor observations where each observation is of length t . The system **100** aims to predict the future blood glucose levels within a specific interval, referred to as a prediction horizon. If n represents the total number of blood glucose readings to be predicted, a target variable will be $Y = \{y_1, y_2, \dots, y_n\}$. Hence, during training, the goal is to learn parameters of the predictive model **120** that can estimate the future blood glucose levels given the input features. Mathematically, a function f is sought that maps the input features X to the predicted blood glucose levels Y :

$$f: \mathbb{R}^{d \times t} \rightarrow \mathbb{R}^n \quad (1)$$

where $\mathbb{R}^{d \times t}$ represents the D -dimensional feature space of length t and $\mathbb{R}^n$ represents the n -dimensional space of the output, for the predicted blood glucose levels.

[0058] In addition to the prediction model **120**, GlySim has an intervention module that allows one to simulate virtual interventions. If the predicted blood glucose level is above a hyperglycemic threshold, i.e. $\max(f(X)) \geq 178$, an intervention can be recommended to prevent the impending hyperglycemia.

[0059] The system **100** can apply the set of multimodal observation data **12** as input to the predictive model **120**, which forecasts future time-dependent values of the factor of interest (e.g., CGM) based on the set of multimodal observation data **12**. In particular, the predictive model **120** generates the sequence of prediction values **14** for the factor of interest across a prediction horizon interval, which may encompass a second time interval (e.g., a few minutes in the future to one or more hours in the future) that follows the first time interval.

[0060] FIG. **3** shows a high-level diagram **300** that illustrates the prediction approach. To this end, the predictive model **120** uses a neural architecture to forecast future blood glucose levels based on multimodal data.

[0061] FIG. **4** shows a neural network architecture of the predictive model **120**. For glucose level prediction, GlySim includes a 3-layer stacked CNN-

LSTM network architecture. One approach to realize such an architecture is to apply a sliding window **410** on the multivariate time series (e.g., the set of multimodal observation data **12**) and slice it to produce sequences of length t for input batches X which can be provided as input to an input layer **420**. A convolutional layer **430**, with tanh activation, applies filters to the input data to extract high-level relevant features. Therefore, if $\mathbf{W}_{conv}$ and $\mathbf{b}_{conv}$ refer to the weights and biases of the convolution layer, the output can be written as:

$$\mathbf{h}_{conv} = \tanh(\mathbf{W}_{conv} * X + \mathbf{b}_{conv}) \quad (2)$$

[0062] A maxpool layer **440** can further consolidate the output of the convolutional layer **430** to prepare the data for input to a Long Short Term Memory (LSTM) stack **450** which can include more than one LSTM layer.

[0063] Each LSTM layer of the LSTM stack **450** processes the sequential data, capturing both short-term and long-term dependencies. It takes the sequence of feature maps produced by the convolutional layer as input. Each LSTM layer maintains a memory cell that can store and update information over time. The LSTM layer uses gates (input, forget, and output) to control the flow of information within the memory cell. Each LSTM layer can learn to remember or forget specific information based on the patterns in the input sequence. By stacking multiple LSTM layers into the LSTM stack **450**, the network can capture more complex temporal dependencies. If o_t is from final output gate and c_t is from final cell state, an output of LSTM layer can be written as:

$$\mathbf{h}_{LSTM} = o_t * \tanh(c_t) \quad (3)$$

[0064] An output layer **460** of the predictive model **120** are configured to generate multi-task outputs, where each task corresponds to a different prediction horizon (i.e., 5-minute, 30-minute, 60-minute).

[0065] FIGS. **5A** and **5B** show alternate configurations of the predictive model **120**. In particular, FIG. **5A** shows a general structure where a quantity of numbers and other parameters may vary based on the problem domain and dataset. FIG. **5B** shows a specific example implementation for CGM prediction.

[0066] During training, a custom loss function can be used to apply varying degrees of error for different regions. In this domain, errors in predicting

glucose levels in hyperglycemic or hypoglycemic regions are more dangerous than those in the normal range.

[0067] It was observed that a major source of error occurs when the predictive model **120** cannot accurately capture the shape of the peaks in the signal. These peaks, however, are quite important because glucose peaks can also represent abnormal health events such as hyperglycemia and hypoglycemia. As such, the custom loss function includes weights (which are essentially penalties) for glucose spikes (i.e., when the true glucose value is outside the normal range). This all happens during training of the predictive model **120**. Results outlined further herein show that this approach improves the CGM prediction/simulation performance significantly.

[0068] Loss Optimization Steps:

[0069] 1. Define thresholds:

$$\text{Lower threshold: } T_{lower} \quad (4.1)$$

$$\text{Upper threshold: } T_{upper} \quad (4.2)$$

[0070] In an example, $T_{lower} = 70 \text{ mg/dl}$ and $T_{upper} = 180 \text{ mg/dl}$.

[0071] 2. Calculate Mean Absolute Error (MAE):

$$MAE = |y_{pred} - y_{true}| \quad (5)$$

[0072] 3. Calculate error in the normal region:

$$T_{lower} \leq y_{true} \leq T_{upper} \quad (6.1)$$

$$normal\_abs\_error = MAE \times 1_{\{T_{lower} \leq y_{true} \leq T_{upper}\}} \quad (6.2)$$

where 1 is an indicator function that is 1 if the condition is true and 0 otherwise.

[0073] 4. Calculate penalties for specific regions:

$$\text{Hypoglycemia region: } y_{true} < T_{lower} \quad (7.1)$$

$$penalty\_lower = 2 \times MAE \times 1_{\{y_{true} < T_{lower}\}} \quad (7.2)$$

$$\text{Hyperglycemia region: } y_{true} > T_{upper} \quad (8.1)$$

$$penalty\_upper = 2 \times MAE \times 1_{\{y_{true} > T_{upper}\}} \quad (8.2)$$

Note that while penalties for errors within fiducial regions (e.g., hyperglycemia or hypoglycemia) are doubled in the examples of Eqs. (7.2) and (8.2), other penalty values may also be assigned which may be fine-tuned to optimize the system.

[0074] 5. Combine the errors and penalties:

$$total\_loss = normal\_abs\_error + penalty\_lower + penalty\_upper \quad (9)$$

B. Intervention Simulation

[0075] Beside forecasting future blood glucose levels with the predictive model, an approach is introduced to simulate interventions and identify an intervention that can lead to preventing dysglycemia (e.g., a hyperglycemic condition). Once the predictive model **120** predicts values above the hyperglycemic threshold, an iterative process is triggered (e.g., at intervention block **130** of FIG. 1) to determine the required reduction in consumed CHO amount from the input data in order to maintain blood glucose levels within the desired range.

[0076] Following generation of the sequence of prediction values, the system **100** evaluates, based on the sequence of prediction values for the factor of interest, whether a predicted maximum value or a predicted minimum value of the sequence of prediction values is outside of the target range.

[0077] In examples where the factor of interest is CGM, a target range may encompass normal CGM readings in view of medical literature to avoid hyper- or hypoglycemia. In some examples, the target range can be modified in view of user attributes such as age, in addition to activities or food intake reported by the user (which may be present within the set of multimodal observation data), e.g., if the user has just eaten a meal, their blood glucose level is not expected to be within “fasting” range. Optionally, the definitions of the target range and/or associated decision processes can factor in uncertainty or how close to “unacceptable” that a predicted value of the sequence of prediction values may be.

[0078] If the predicted maximum value or a predicted minimum value of the sequence of prediction values is outside of the target range, then the system **100** can initiate an intervention process that identifies an intervention action that can be applied with respect to the plurality of parameters in order to keep or otherwise cause the factor of interest to be within the target range.

[0079] For example, the consumed *carbohydrate* (or any other parameter effects blood glucose) amount can be manipulated in input feature X in an iterative process to identify the optimal reduction in CHO to achieve a predicted CGM value below hyperglycemic threshold. Algorithm 1 is shown below for further illustration of an iterative process for determining an intervention action.

[0080] With this intervention mechanism in place, the aim is to optimize the consumed CHO amount and maintain blood glucose levels within the desired range for effective glucose management.

Algorithm 1: CHO reduction-based intervention

Inputs:

multimodal features, $X' = \{x^{(1)}, x^{(2)}, \dots, x^{(d-1)}\}$

projected CHO intake, C

hyperglycemic threshold, BG_{max}

predictive model, $f(\cdot)$

Output: CHO reduction, C_r

if $\max(f(X \oplus C)) \geq BG_{max}$ **then**

while $\max(f(X \oplus C)) \geq BG_{max}$ **do**

$C_r += 1$

end while

return C_r

else

$C_r = 0$

end if

[0081] The system **100** can determine, over one or more iterations, an intervention action to be applied at an intervention time step within the second time interval with respect to the plurality of parameters, where the intervention action is predicted to result in the *intervention sequence of prediction values* **24** for the factor of interest having a predicted maximum value and a predicted minimum value that are within a target range. To achieve this, the system **100** can modify a value of a parameter that can be expected to affect the value of the factor of interest, such as consumed carbohydrate amount or insulin dosage, and use the predictive model **120** to simulate what would happen to the factor of interest under the modified value.

1. *Generating an intervention set of multimodal observation data*

[0082] To determine the intervention action, the system **100** can generate, for an iteration of the one or more iterations, the intervention set of multimodal observation data **22** that includes a modified value of a parameter of the plurality of parameters for the intervention time step within the second time interval, the modified value corresponding to a potential intervention action.

[0083] For example, if the factor of interest is a blood glucose level where the sequence of prediction values is too high (e.g., hyperglycemia), an intervention set of multimodal observation data **22** can include a reduced carbohydrate amount relative to the (original) carbohydrate amount from the (original) set of multimodal observation data **12** or a carbohydrate amount from a previous iteration of the intervention set of multimodal observation data **22**. Alternatively, for the hyperglycemia example, an intervention set of multimodal observation data **22** can include an increased amount of basal or bolus insulin.

2. *Simulating to obtain intervention sequence of prediction values*

[0084] The system **100** can then apply the *intervention set of multimodal observation data 22* as input to the predictive model **120**, which in some examples can correspond with an intervention time step (which may be after the first interval associated with the (actual) set of multimodal observation data **12**). The predictive model **120** can generate the *intervention sequence of prediction values 24* for the factor of interest which would occur under the modified parameter values expressed within the intervention set of multimodal observation data **22**. Corresponding to the hyperglycemia example where the intervention set of multimodal observation data **22** includes a reduced carbohydrate amount, the predictive model **120** simulates blood glucose levels that would result from the reduced carbohydrate amount specified within the intervention set of multimodal observation data **22**.

3. *Iterative searching for optimal intervention action*

[0085] In the hyperglycemia example, the blood glucose levels would be expected to reduce as the carbohydrate amount is reduced. However, the correct *reduction amount* would need to be found to ensure that the intervention action is enough to avoid hyperglycemia.

[0086] As such, the system **100** can evaluate, for the iteration of the one or more iterations and based on the intervention sequence of prediction values **24** for the factor of interest within the prediction horizon interval, whether the predicted maximum value or the predicted minimum value of the intervention sequence of prediction values **24** are outside of the target range or are otherwise sub-optimal. If so, then the system **100** continues searching for an appropriate intervention action with further modification to the intervention set of multimodal observation data **22** over further iterations. If not, then the system **100** can end the process and report the parameter value associated with the intervention set of multimodal observation data **22** as a potential intervention action.

User Interface

[0087] Referring to FIGS. **5** and **6A-6D**, the system **100** can generate and display, at the display device **230** (FIG. **2**), a graphical representation showing one or more of: the set of multimodal observation data **12**, the sequence of prediction values **14**, the intervention set of multimodal observation data **22** associated with an intervention action, and/or the intervention sequence of prediction values **24** associated with an intervention action. A diagram **500** in FIG. **5** shows a “dashboard” of the graphical user interface **140** for implementation of GlySim.

[0088] A user can enter information that can be used to construct the set of multimodal observation data **12**, such as meal times and statistics (e.g., carb amounts), insulin doses and times, exercise, work, sleep, and other items. The user can also control simulation factors, and can initiate simulation of a factor of interest based on inputs they provide. The user can, in some examples, specify potential changes to parameters that may be used to construct the intervention set of multimodal observation data **22** for simulation (e.g., “what would happen to my blood glucose levels if I were to eat this cupcake?”).

[0089] FIGS. **7A-7D** show an example where GlySim detects hyperglycemia beforehand, detects increasing prior BGL is forcing predicted BGL up, and reduces prior BGL iteratively until simulation shows normoglycemia. FIG. **7A** shows example values of a parameter of the set of multimodal observation data **12** which corresponds with an example sequence of prediction values **14** in FIG. **7B**. FIG. **7C** shows example values of a parameter of the intervention set of multimodal

observation data **22** which corresponds with an example intervention sequence of prediction values **24** in FIG. **7D**.

Experimental Implementation

A. Dataset

[0090] The OhioT1DM was selected to demonstrate the results of the methods outlined herein. OhioT1DM was collected over an eight-week long clinical study of 12 deidentified T1 diabetes patients. Participants wore Medtronic 530G/630G insulin pumps and Medtronic Enlite CGM that transmits blood glucose level every 5 minutes. Their physiological data (acceleration, skin response etc.) was recorded on either Basis Peak fitness band or Empatica Embrace while their self-reported CHO intake, work, exercise intensity and sleep quality were recorded on smartphones. Of these eight weeks, roughly 44 and 12 days were allocated for train and test respectively. Missing values were imputed with interpolation or extrapolation. However, subjects 540 and 567 were excluded from analysis due to not having entries for CHO, work, sleep, and exercise sessions.

[0091] For implementing the loss function in Eqs. (4.1)-(9), CGM data was classified based on two common thresholds for hyperglycemia and hypoglycemia:

[0092] If $CGM < 70$, then class label = 0.

[0093] If $70 \leq CGM \leq 180$, then class label = 1.

[0094] If $CGM > 180$, then class label = 2.

[0095] This is crucial for two reasons: (1) the categorization allows the model to handle regions differently; and (2) the custom loss function in Eqs. (4.1)-(9) may apply heavier penalties for sensitive areas.

B. Model Development

[0096] After smoothing the CGM stream with Kalman filter, 180-minute long multimodal time-series windows were produced with 175-minute overlap and fed as input to a network with one 1-D CNN layer and two LSTM layers, all equipped with tanh activation and followed by a 50% dropout. The CNN layer had 64 filters, a kernel size of 3, l_2 regularizer followed by a maxpool. The two LSTM layers had 50 units each. The output specific layers were for 60-minute, 30-minute and immediate

next sample prediction each with linear activation. Adam optimizer was used with a learning rate of 0.01 and a decay of 0.001. Grad-CAM was also used in the intervention phase to identify the responsible factor(s) behind hyperglycemic/hypoglycemic events.

C. Loss Function

[0097] FIG. **8A** shows results from a version of the predictive model **120** that was implemented without categorization and separate penalties for different regions in its loss function, and FIG. **8B** shows results from a version of the predictive model **120** that was implemented with the loss function discussed herein with respect to Eqs. (4.1)-(9). The loss function in Eqs. (4.1)-(9) is theorized to modify the loss function to improve regression accuracy (or reduce regression error). Results for FIG. **8A** show an RMSE of 28.13, an MSE of 791.51, an MAE of 19.06, and an MAPE of 0.13%. Results for FIG. **8B** show an RMSE of 26.87, an MSE of 721.99, an MAE of 18.63, and an MAPE of 0.13%.

D. Validation

[0098] FIGS. **9A-9C** are a series of graphical representations showing ground truth CGM values and predicted CGM values for three patients from the OhioT1DM Dataset, which are also represented in Table **1** below.

[0099] FIG. **9A** shows ground truth and predicted CGM values for a first patient ("patient ID 570"), resulting in an RMSE of 19.94, an MSE of 397.61, an MAE of 13.52, and an MAPE of 0.07%.

[00100] FIG. **9B** shows ground truth and predicted CGM values for a second patient ("patient ID 588"), resulting in an RMSE of 22.00, an MSE of 483.99, an MAE of 15.63, and an MAPE of 0.09%.

[00101] FIG. **9C** shows ground truth and predicted CGM values for a third patient ("patient ID 563"), resulting in an RMSE of 23.59, an MSE of 556.25, an MAE of 15.70, and an MAPE of 0.10%.

[00102] Table **1** shows comparison of RMSE between GlySim and three other state-of-the-art methods for the three patients discussed above with reference to FIGS. **9A-9C**.

TABLE 1: Error Comparison

Patient ID	Bi-LSTM	Martinsson, et al. 2018	Marking, et al. 2020	GlySim
570	25.48	28.6	25.74	19.94
588	30.17	33.1	30.45	22.00
563	29.11	29.9	30.42	23.59
Mean RMSE	28.25	30.53	28.87	21.84

Computer-implemented System

[00103] FIG. 2 is a schematic block diagram of an example computing device **200** that may be used with one or more embodiments described herein, e.g., as a component implementing GlySim.

[00104] Device **200** comprises one or more network interfaces **210** (e.g., wired, wireless, PLC, etc.), at least one processor **220**, and a memory **240** interconnected by a system bus **250**, as well as a power supply **260** (e.g., battery, plug-in, etc.). Computing device **200** can also include a display device **230** for displaying aspects of the GUI **140** including graphical representations such as those shown in FIGS. 6 and 7A-7D.

[00105] Network interface(s) **210** include the mechanical, electrical, and signaling circuitry for communicating data over the communication links coupled to a communication network. Network interfaces **210** are configured to transmit and/or receive data using a variety of different communication protocols. As illustrated, the box representing network interfaces **210** is shown for simplicity, and it is appreciated that such interfaces may represent different types of network connections such as wireless and wired (physical) connections. Network interfaces **210** are shown separately from power supply **260**, however it is appreciated that the interfaces that support PLC protocols may communicate through power supply **260** and/or may be an integral component coupled to power supply **260**.

[00106] Memory **240** includes a plurality of storage locations that are addressable by processor **220** and network interfaces **210** for storing software programs and data structures associated with the embodiments described herein. In some embodiments, device **200** may have limited memory or no memory (e.g., no memory for storage other than for programs/processes operating on the device and associated caches). Memory **240** can include instructions executable by the

processor **220** that, when executed by the processor **220**, cause the processor **220** to implement aspects of GlySim and the methods outlined herein.

[00107] Processor **220** comprises hardware elements or logic adapted to execute the software programs (e.g., instructions) and manipulate data structures **245**. An operating system **242**, portions of which are typically resident in memory **240** and executed by the processor, functionally organizes device **200** by, inter alia, invoking operations in support of software processes and/or services executing on the device. These software processes and/or services may include GlySim processes/services **290**, which can include aspects of methods and/or implementations of various modules described herein, including those shown and discussed above with respect to FIGS. 1 and 3-7D.. Note that while GlySim processes/services **290** is illustrated in centralized memory **240**, alternative embodiments provide for the process to be operated within the network interfaces **210**, such as a component of a MAC layer, and/or as part of a distributed computing network environment.

[00108] It will be apparent to those skilled in the art that other processor and memory types, including various computer-readable media, may be used to store and execute program instructions pertaining to the techniques described herein. Also, while the description illustrates various processes, it is expressly contemplated that various processes may be embodied as modules or engines configured to operate in accordance with the techniques herein (e.g., according to the functionality of a similar process). In this context, the term module and engine may be interchangeable. In general, the term module or engine refers to model or an organization of interrelated software components/functions. Further, while the GlySim processes/services **290** is shown as a standalone process, those skilled in the art will appreciate that this process may be executed as a routine or module within other processes.

[00109] It should be understood from the foregoing that, while particular embodiments have been illustrated and described, various modifications can be made thereto without departing from the spirit and scope of the invention as will be apparent to those skilled in the art. Such changes and modifications are within the scope and teachings of this invention as defined in the claims appended hereto.

CLAIMS

What is claimed is:

1. A system comprising:
 - a processor in communication with a memory and one or more sensor devices, the memory including instructions executable by the processor to:
 - access, by a processor in communication with a memory and one or more sensor devices, a set of multimodal observation data for a first time interval, the set of multimodal observation data including values of a plurality of parameters that affect values of a factor of interest and the set of multimodal observation data including time-dependent values for the factor of interest;
 - generate, by application of the set of multimodal observation data as input to a prediction model, a sequence of prediction values for the factor of interest across a prediction horizon interval, the prediction horizon interval including a second time interval following the first time interval associated with the set of multimodal observation data; and
 - determine, over one or more iterations, an intervention action to be applied at an intervention time step within the second time interval with respect to the plurality of parameters that is predicted to result in an intervention sequence of prediction values for the factor of interest having a predicted maximum value and a predicted minimum value that are within a target range.
2. The system of claim 1, the factor of interest including a blood glucose level, and the time-dependent values of the factor of interest for the first time interval being measurable by a blood glucose sensor of the one or more sensor devices.

3. The system of claim 1, the plurality of parameters including one or more of a consumed carbohydrate amount, a basal insulin dosage or a bolus insulin dosage.
4. The system of claim 1, the set of multimodal observation data associated with the first time interval including one or more physiological signals.
5. The system of claim 1, the memory further including instructions executable by the processor to:
 - evaluate, based on the sequence of prediction values for the factor of interest associated with the set of multimodal observation data, whether a predicted maximum value or a predicted minimum value of the sequence of prediction values is outside of the target range.
6. The system of claim 1, further comprising a display device, the memory further including instructions executable by the processor to:
 - display a graphical representation at the display device that illustrates the set of multimodal observation data for the first time interval and the sequence of prediction values for the factor of interest across the prediction horizon interval.
7. The system of claim 1, the memory further including instructions executable by the processor to:
 - generate, for an iteration of the one or more iterations, an intervention set of multimodal observation data that includes a modified value of a parameter of the plurality of parameters for the intervention time step within the second time interval, the modified value corresponding to the intervention action;
 - generate, for the iteration of the one or more iterations and by application of the intervention set of multimodal observation data as input to the prediction model, the intervention sequence of prediction values for the factor of interest following the intervention time step; and
 - evaluate, for the iteration of the one or more iterations and based on the sequence of prediction values for the factor of interest within the

prediction horizon interval, whether the predicted maximum value or the predicted minimum value of the intervention sequence of prediction values is outside of the target range.

8. The system of claim 7, further comprising a display device, the memory further including instructions executable by the processor to:
 - display a graphical representation at the display device that illustrates the intervention set of multimodal observation data and the intervention sequence of prediction values for the factor of interest.
9. The system of claim 1, the memory further including instructions executable by the processor to:
 - apply a sliding window to the set of multimodal observation data resulting in an input feature sequence;
 - extract, by a convolutional layer of the prediction model, a high-level feature sequence based on the input feature sequence; and
 - generate, by a long short term memory layer of the prediction model, the sequence of prediction values for the factor of interest across the prediction horizon interval based on the high-level feature sequence.
10. The system of claim 9, the memory further including instructions executable by the processor to:
 - modify the input feature sequence based on an intervention set of multimodal observation data;
 - extract, by the convolutional layer of the prediction model, an intervention high-level feature sequence based on the input feature sequence; and
 - generate, by the long short term memory layer of the prediction model, the intervention sequence of prediction values for the factor of interest based on the intervention high-level feature sequence.
11. The system of claim 1, the memory further including instructions executable by the processor to:
 - optimize a loss function that more heavily penalizes errors associated with the sequence of prediction values or a ground truth sequence of values

when the sequence of prediction values or the ground truth sequence of values are within a specified region.

12. A method, comprising:
 - accessing, by a processor in communication with a memory and one or more sensor devices, a set of multimodal observation data for a first time interval, the set of multimodal observation data including values of a plurality of parameters that affect values of a factor of interest and the set of multimodal observation data including time-dependent values for the factor of interest;
 - generating, by application of the set of multimodal observation data as input to a prediction model, a sequence of prediction values for the factor of interest across a prediction horizon interval, the prediction horizon interval including a second time interval following the first time interval associated with the set of multimodal observation data; and
 - determining, over one or more iterations, an intervention action to be applied at an intervention time step within the second time interval with respect to the plurality of parameters that is predicted to result in an intervention sequence of prediction values for the factor of interest having a predicted maximum value and a predicted minimum value that are within a target range.
13. The method of claim 12, further comprising:
 - evaluating, based on the sequence of prediction values for the factor of interest associated with the set of multimodal observation data, whether a predicted maximum value or a predicted minimum value of the sequence of prediction values is outside of the target range.
14. The method of claim 12, further comprising:
 - displaying a graphical representation at a display device that illustrates the set of multimodal observation data for the first time interval and the sequence of prediction values for the factor of interest across the prediction horizon interval.

15. The method of claim 12, further comprising:
 - generating, for an iteration of the one or more iterations, an intervention set of multimodal observation data that includes a modified value of a parameter of the plurality of parameters for the intervention time step within the second time interval, the modified value corresponding to the intervention action;
 - generating, for the iteration of the one or more iterations and by application of the intervention set of multimodal observation data as input to the prediction model, the intervention sequence of prediction values for the factor of interest following the intervention time step; and
 - evaluating, for the iteration of the one or more iterations and based on the sequence of prediction values for the factor of interest within the prediction horizon interval, whether the predicted maximum value or the predicted minimum value of the intervention sequence of prediction values is outside of the target range.
16. The method of claim 15, further comprising:
 - displaying a graphical representation at a display device that illustrates the intervention set of multimodal observation data and the intervention sequence of prediction values for the factor of interest.
17. The method of claim 12, the factor of interest including a blood glucose level, and the time-dependent values of the factor of interest for the first time interval being measurable by a blood glucose sensor of the one or more sensor devices.
18. The method of claim 12, the plurality of parameters including a consumed carbohydrate amount.
19. The method of claim 12, further comprising:
 - applying a sliding window to the set of multimodal observation data resulting in an input feature sequence;
 - extracting, by a convolutional layer of the prediction model, a high-level feature sequence based on the input feature sequence; and

generating, by a long short term memory layer of the prediction model, the sequence of prediction values for the factor of interest across the prediction horizon interval based on the high-level feature sequence.

20. The method of claim 19, further comprising:
- modifying the input feature sequence based on an intervention set of multimodal observation data;
 - extracting, by the convolutional layer of the prediction model, an intervention high-level feature sequence based on the input feature sequence; and
 - generating, by the long short term memory layer of the prediction model, the intervention sequence of prediction values for the factor of interest based on the intervention high-level feature sequence.




FIG. 1




FIG. 2




FIG. 3




FIG. 4





FIG. 5A






FIG. 6




FIG. 7A




FIG. 7B




FIG. 7C





FIG. 7D






FIG. 8B




FIG. 9A




FIG. 9B




FIG. 9C