SYSTEMS, METHODS, AND DEVICES FOR BIOPHYSICAL MODELING AND RESPONSE PREDICTION

CROSS-REFERENCE

[0001] This application claims priority to U.S. Provisional Patent Application No. 62/773,117, filed on November 29, 2018, U.S. Provisional Patent Application No. 62/773,125, filed on November 29, 2018, U.S. Provisional Patent Application No. 62/773,134, filed on November 29, 2018, each of which is incorporated herein by reference in its entirety.

BACKGROUND

[0002] Many biological systems, including the human body, can function with a high degree of complexity. Further, while a single species can have an overall general likeness, there can be significant variability between individuals. Consequently, it can be difficult to understand, let alone predict various biological responses on an individual basis.

[0003] Understanding human biological responses on an individual basis can provide various health and quality-of-life benefits. Such an understanding can enable an individual to make better choices to improve their health. When such choices are made by a population, the overall health of society can benefit. In addition, such an understanding can empower an individual to better alter their lifestyle in the pursuit of personal goals.

[0004] One human biological response of increasing interest is nutrition, and blood glucose levels resulting from eating in particular. Failure to maintain blood glucose levels in acceptable levels over time may result in adverse consequences, including pre-diabetes or Type 2 diabetes. However, individuals can vary in blood glucose response, diet, behavior, and numerous other factors. Accordingly, conventional models (e.g., linear blood glucose response models) can be inadequate understanding an individual's personalized blood glucose response.

SUMMARY

[0005] The present disclosure provides systems and methods that can acquire sensor and other data that records subject actions and that utilize reinforcement learning to predict a subject response and the modify the subject's response to achieve a goal. The prediction can be a prediction of a biophysical response and/or behavioral response. Embodiments can utilize custom variational encoding to model subject actions and responses. In some embodiments, the systems described herein can generate a recommendation for a subject based on a reward function and the subject's historical actions.

[0006] In one aspect, the present disclosure provides a computer-implemented method for training a reinforcement learning algorithm.

The method can comprise: (a) providing a reaction model comprising a personalized learned body model configured to generate a simulated biophysical or behavioral response of a subject responsive to a plurality of inputs; (b) until a convergence condition is achieved, iteratively: (i) generating, using the reinforcement learning algorithm, a recommendation comprising a recommended meal or physical activity, (ii) processing the recommendation using the reaction model to generate a predicted response of the subject to following the recommendation, wherein the predicted response comprises a biophysical or behavioral response, (iii) applying a first reward function to the predicted response to generate a first reward value indicative of a benefit or detriment of the predicted response toward a health measurement of interest, and (iv) training the reinforcement learning algorithm in a first stage using the first reward value. The method can further comprise (c) providing the recommendation to the subject; measuring the health measurement oof interest of the subject responsive to following the recommendation; and € applying a second reward function to the measured health measurement of interest to generate a second reward value and (f) training the reinforcement learning algorithm in a second stage using the second reward value. In some embodiments, the method comprises using the glucose response of the subject to train the biophysical reaction model. In some embodiments, the method comprises encoding the glucose response of the subject into a low-dimension latent space for providing to the biophysical reaction model and to the second reward function. In some embodiments, the first reward function is the same as the second reward function. In some embodiments, the biophysical reaction model includes at least one body model configured to generate a simulated biophysical response of the subject in response to a plurality of inputs. In some embodiments, generating the predicted glucose response of the subject comprises: applying the glucose response of the subject to a predictor trained to infer a future glucose response; applying the recommendation to an adherence model configured to evaluate how closely the subjects will follow the recommendation; and selectively applying outputs of the predictor and the adherence model as the plurality of inputs to the body model. In some embodiments, generating the predicted glucose response further comprises applying the simulated biophysical response to an autoencoder and generative adversarial network to generate the predicted glucose response. In some embodiments, the convergence condition is based at least in part on the magnitude of the first reward.

[0006a] In a second aspect, the present disclosure provides a system comprising one or more computer processors and one or more storage devices having instructions stored thereon that are operable, when executed by the one or more computers, to cause the one or more computer

processors to perform operations comprising: (a) providing a reaction model comprising a personalized learned body model configured to predict a biophysical or behavioral response of a subject responsive to a plurality of inputs; (b) until a convergence condition is achieved, iteratively: (i) generating, using the reinforcement learning algorithm, a recommendation comprising a recommended meal or physical activity, (ii) processing the recommendation using the reaction model to generate a predicted response of the subject to following the recommendation, wherein the predicted response comprises a biophysical or behavioral response, (iii) applying a first reward function to the predicted response to generate a first reward value indicative of a benefit or detriment of the predicted response toward a health measurement of interest; and (iv) training the reinforcement learning algorithm in a first stage using the first reward value; (c) providing the recommendation to the subject; (d) measuring the health measurement of interest of the subject responsive to following the recommendation; (e) applying a second reward function to the measured health measurement of interest to generate a second reward value; and (f) training the reinforcement learning algorithm in a second stage using the second reward value

[0006b] In a third aspect, the present disclosure provides a non-transitory computer storage medium having instructions stored thereon that are operable, when executed by one or more computer processors, to cause the one or more computer processors to perform operations comprising: (a) providing a reaction model comprising a personalized learned body model configured to predict a biophysical or behavioral response of a subject responsive to a plurality of inputs; (b) until a convergence condition is achieved, iteratively: (i) generating, using the reinforcement learning algorithm, a recommendation comprising a recommended meal or physical activity, (ii) processing the recommendation using the reaction model to generate a predicted response of the subject to following the recommendation, wherein the predicted response comprises a biophysical or behavioral response, (iii) applying a first reward function to the predicted response to generate a first reward value indicative of a benefit or detriment of the predicted response toward a health measurement of interest; and (iv) training the reinforcement learning algorithm in a first stage using the first reward value; (c) providing the recommendation to the subject; (d) measuring the health measurement of interest of the subject responsive to following the recommendation; (e) applying a second reward function to the measured health measurement of interest to generate a second reward value; and (f) training the reinforcement learning algorithm in a second stage using the second reward value.

[0007] In another example, the present disclosure provides a method that can comprise: training each of a plurality of different autoencoder (AE) temporal convolutional neural

networks (CNNs) on historical time series data from one of a plurality of different data sources, wherein the plurality of different data sources comprises a continuous glucose monitor, a heart rate monitor, and a source of food data; generating a plurality of seed values using the plurality of AE temporal CNNs in response to current data from the plurality of different data sources; configuring each of a plurality of different CNN encoders with one of the plurality seed values

from a corresponding AE temporal CNN; applying past the historical times series data from the corresponding data sources to the temporal CNN encoders to generate encoded data values; and applying the encoded data values to a forecast configured to generate predicted data values corresponding to one of the data sources. In some embodiments, training the plurality of AE temporal CNNs includes: training a first AE temporal CNN with data from a first type sensor that reads a first biophysical response, and training a second AE temporal CNN with data from a second type sensor that reads a second biophysical response; and applying the encoded data values to generate predicted first biophysical responses.

[0008] In another example, the present disclosure provides a method that can comprise: receiving reference values over time from a first biophysical sensor that represent at least one biophysical response of a subject, the time including a first time period followed by a second time period; receiving first data values from a second biophysical sensor; inferring first predicted values for the at least one biophysical response with a first subject model using at least first data from the second time period; inferring second predicted values for the at least one biophysical response with a second subject model using at least first data from the first time period; and comparing the first and second predicted values to the reference values to determine the accuracy of the subject models. In some embodiments, inferring second predicted values includes: inferring predicted first data values for the second time period from first data values of the first time period, and applying the first predicted data values for the second time period to the second subject model. In some embodiments, the method further comprises: receiving second data values; inferring first predicted values for the at least one biophysical response further includes using second data from the second time period; and inferring second predicted values for the at least one biophysical response further includes using second data from the first time period. In some embodiments, the method further comprises: in response to at least the reference values and first data values from the first time period, adjusting parameters in the second subject model. In some embodiments, the method further comprises: applying the first predicted values for the at least one biophysical response to a first generative adversarial network to generate first adjusted predicted values; applying the second predicted values for the at least one biophysical response to a second generative adversarial network to generate second adjusted predicted values; and comparing the first and second adjusted predicted values to the reference values to determine the accuracy of the generative adversarial networks. In some embodiments, the method further comprises: in response to comparing the first and second predicted values to the reference values updating at least the first and second body models. In another example, the

present disclosure provides a system that can comprise: a first data prediction model configured to

generate predicted first data values for a second time period from first data values for a first time period which precedes the second time period; a second data prediction model configured to generate predicted second data values for the second time period from second data values for the first time period; a first subject model comprising at least in part an artificial neural network (ANN) configured to infer a first predicted biophysical response from first and second data values from the second time period; a second subject model having a same structure as the first body model configured to infer a second predicted biophysical response from predicted first and second data values for the second time period; and a compare system configured to compare a reference biophysical response for the second time period to the first and second predicted biophysical responses. In some embodiments, the first and second data prediction models comprise long short-term memory networks. In some embodiments, the system further comprises: a third data prediction model configured to generate a predicted reference biophysical response for the second time period from the reference biophysical response for the first time period. In some embodiments, the system further comprises: a parameter estimator configured to update at least the second subject model in response to a reference biophysical response, and the first and second data values from the first time period. In some embodiments, the system further comprises: a feedback generator configured to selectively adjust any of the first data prediction model, a second data prediction model, the first subject model, the second subject model in response to comparison results from the compare system.

[0009] In another example, the present disclosure provides a method that can comprise: during a first time interval, training a biophysical model comprising an artificial neural network (ANN) to generate a simulated biophysical response of a subject from at least first sensor data and second sensor data, the first sensor data comprising continuous glucose monitoring data and the second sensor data comprising heart rate monitoring data, which training comprises estimating personalized time-varying parameters of the biophysical model; during a second time interval, generating the simulated biophysical response of the subject in real time using the trained biophysical model, the real-time data including the second sensor data but not including the first sensor data. In some embodiments, training the biophysical model includes: applying at least the first sensor data to the biophysical model, and updating the biophysical model using a parameter estimator that evaluates the simulated first biophysical response with respect to a subject's actual first biophysical response. In some embodiments, training the biophysical model includes applying the first sensor data and the second data to the first biophysical model, the second data being different than the first biophysical response. In some embodiments, the second data comprises data logged by the subject.

[0010] In another example, the present disclosure provides a system that can comprise: a first biophysical model comprising an artificial neural network (ANN) configured to derive network parameters in response to training with data from at least one type of sensor to output a simulated first type biophysical response; a parameter estimator configured to update parameters of the first biophysical model in response to the simulated first type biophysical response and an actual first type biophysical response of a subject; and a second biophysical model comprising an ANN, configured with the network parameters, operable to infer predicted first type biophysical responses in real-time for the subject in response to sensor data received from the at least one type of sensor in real-time. In some embodiments, the actual first type biophysical response is recorded with a first type sensor and the at least one type of sensor is different than the first type sensor. In some embodiments, the at least one type of sensor records a second type biophysical response different from the first type biophysical response. In some embodiments, the actual first type biophysical response is recorded with a first type sensor and the at least one type of sensor is different than the first type sensor. In some embodiments: the first biophysical model derives the network parameters in response to training with data from at least one type of sensor and data from a second source; and the second biophysical model infers predicted first type biophysical responses in response to sensor data received from the at least one type of sensor in real-time and data from the second source received in real-time. In some embodiments, the second source comprises data logged personally by the subject. In another example, the present disclosure provides a method that can comprise: receiving time series data from a plurality of different sources that each record data of a different type for at least one subject that performs actions, the different sources including a glucose sensor that records a glucose response of the at least one subject; executing unsupervised learning on the time series data with at least one encoding artificial neural network (ANN) to produce encoded values in a resulting latent space having a predetermined distance from one another; selecting orthogonal values based on the latent space; decoding the orthogonal values with an ANN having a corresponding decoding structure to the encoding ANN to generate decoded values; and mapping the decoded values to subject actions. In some embodiments, executing unsupervised learning includes autoencoding the time series data. In some embodiments, autoencoding the time series data includes autoencoding with a temporal convolutional neural network (NN) variational autoencoder. In some embodiments, the method further comprises filtering the decoded values based on relevance criteria for a particular subject. In some embodiments, the different sources further include a second type sensor that records a second type biophysical response of the at least one subject. In some embodiments, at least one subject action is selected from physical activities of the subject and the ingestion of

food by the subject. In some embodiments, at least one of the different sources is selected from the group of accelerometer data, calendar data of the subject, and sleep state of the subject.

[0011] In another example, the present disclosure provides a system that can comprise: an encoder configured to encode time series data values into encoded values in a latent space having a predetermined metric distance from one another, the time series data being from a plurality of different data sources that record features of at least one subject, at least one time series data being for a first type biophysical response of the at least one subject; a value selector module configured to determine orthogonal values from the encoded values; an decoder having a decoding structure corresponding to the encoder and configured to generate decoded values from the orthogonal values; and an action mapping module configured to map the decoded values to actions of the at least one subject. In some embodiments, the autoencoder comprises a temporal convolutional NN variational autoencoder. In some embodiments, the system further comprises a filtering module configured to selectively discard some of the decoded values based on relevance criteria for a particular subject. In some embodiments, at least another of the time series data is for a second type biophysical response of the at least one subject. In some embodiments, the data sources include a continuous glucose meter, heart rate monitor and food data logged by the at least one subject.

[0012] In another example, the present disclosure provides a method that can comprise: creating a data object in a system memory of a computing system; copying data into the data object; by execution of a decorator function, transforming the data object into a data processing object having an egress messaging function; processing the data of the data processing object with one of a plurality of different machine learning processes; and upon completing the processing of the data, returning a processing result and executing the egress messaging function. In some embodiments, the plurality of different processes are asynchronous processes, and wherein the method further comprises, upon receiving a message for the egress messaging function: creating a next data object in the system memory; copying next data into data object; by execution of the decorator function, transforming the next data object into a next data processing object having the egress messaging function; and processing data of the next data processing object with one of the machine learning processes. In some embodiments, the processing result comprises a dictionary object that maps keys to values. In some embodiments, the data comprises input data to an artificial neural network (ANN) for learning operations on the ANN. In some embodiments, the data comprises input data to an artificial neural network (ANN) for inference operations on the ANN.

[0013] In another example, the present disclosure provides a system that can comprise: a data store configured to store data for processing; system memory; a multiprocessing module configured to execute a plurality of machine learning processes in parallel; and a data object decorating function comprising instructions executable by the multiprocessing module and configured to: create a data object in the system memory, copy data into the data object from the data store, transform the data object into a data processing object having an egress messaging function, and instantiate one of the machine learning processes to process the data of the data processing object and return processing results and execute the messaging function to return a message. In some embodiments, the multiprocessing module is resident on a server. In some embodiments, the multiprocessing module is distributed over a plurality of servers. In some embodiments, the machine learning processes include an artificial neural network (ANN). In some embodiments, the ANN is selected from the group consisting of autoencoders (AEs), generative adversarial networks (GANs), long short-term memory networks (LSTMs), convolutional neural networks (CNNs), and reinforcement learning (RL) algorithms.

[0014] In another example, the present disclosure provides a method that can comprise: creating a biophysical model with at least one machine learning architecture to predict a first biophysical response, wherein the biophysical model has been trained with at least primary sensor data and secondary sensor data, the primary sensor data capturing a first biophysical response, the secondary sensor data capturing a second biophysical response; in response to at least the secondary sensor data and not the primary sensor data, predicting a first biophysical response of the subject with the biophysical model; determining if the predicted first biophysical response is outside of predetermined limits; and if the predicted first biophysical response is outside of predetermined limits, transmitting at least one recommendation to the subject, the at least one recommendation selected to adjust the subject's actual biophysical response to be within the predetermined limits. In some embodiments, the method further comprises setting the predetermined limits according to the subject's health status. In some embodiments, the method further comprises setting the predetermined limits according to a subject's health goals. In some embodiments, the first biophysical response includes a glucose response of the subject. In some embodiments, the second biophysical response includes a heart rate of the subject. In some embodiments, the biophysical model is also trained with data logged by the subject. In some embodiments, the primary sensor data is generated from a continuous glucose monitor; the secondary sensor data is generated from a heart rate monitor; and the data logged by the subject is food eaten by the subject. In some embodiments, the at least one recommendation is selected from a physical activity recommendation and a food recommendation. In some embodiments, the

biophysical model comprises an artificial neural network configured as an autoencoder. In some embodiments, the biophysical model comprises an artificial neural network configured as at least a long short-term memory (LSTM) configured to predict a first biophysical response from a recommendation. In some embodiments, the biophysical model comprises an artificial neural network configured as at least one temporal convolutional neural network configured to predict a first biophysical response of the subject. In some embodiments, the at least one recommendation is selected from a recommendation set including canonical actions derived by autoencoding heterogenous sensor data. In some embodiments, the method further comprises, if the predicted first biophysical response is not outside of predetermined limits, transmitting a predetermined message. In some embodiments, the predetermined message is selected from the group of: an encouragement message and a reward. In some embodiments, the method further comprises displaying the at least one recommendation on a subject device. In some embodiments, the method further comprises capturing at least the secondary sensor data with an application executable on a subject device. In some embodiments, the method further comprises capturing the primary data with the application.

[0015] In another example, the present disclosure provides a method that can comprise: training a glucose regulation model having at least one first parameter to predict glucose levels in response to at least food source data; in response to information on a subject, substituting the at least one first parameter with at least one personalized parameter in the glucose regulation model to create a personalized glucose regulation model; and applying food source data from the subject to the personalized glucose regulation model to predict a glucose level of the subject. In some embodiments, the glucose regulation model includes at least one neural network. In some embodiments, the glucose regulation model includes at least one statistical model selected form the group consisting of: a long short-term memory neural network and recurrent neural network. In some embodiments, the glucose regulation model includes at least one neural network trained with data of a predetermined population. In some embodiments, the at least one first parameter comprises an insulin resistance parameter. In some embodiments, the glucose regulation model includes at least one glucose model selected from the group consisting of: a differential equation model of glucose regulation and a glucose model comprising a set of coupled equations. In some embodiments, the at least one differential equation model of glucose regulation includes a food source function. In some embodiments, the method further comprises: training the food source function with at least training data selected from the group consisting of: glycemic responses of a population to predetermined foods, and glycemic responses calculated from data for predetermined foods. In some embodiments, the method further comprises generating the

personalized parameters of the subject by recording a glucose response of the subject with a glucose meter. In some embodiments, the method further comprises generating the personalized parameters of the subject by classifying the subject into a demographic equivalent group based on characteristic data of the subject.

[0016] In another example, the present disclosure provides a system that can comprise: a computing system comprising a glucose prediction model comprising at least one model parameter operable to predict glucose levels in response to at least food source data; a model parameter input configured to receive at least one personalized parameter as the at least one model parameter, the at least one personalized parameter generated in response to data of a subject; and a food source data input configured to apply food source data to the glucose prediction model with the at least one personalized parameter to predict a glucose level of the subject. In some embodiments, the glucose regulation model comprises a neural network. In some embodiments, the glucose regulation model comprises a statistical model selected from the group consisting of: a long short-term memory neural network and recurrent neural network. In some embodiments, the glucose regulation model comprises at least at least one neural network trained with data from a predetermined population. In some embodiments, the at least one model parameter includes an insulin resistance parameter. In some embodiments, the glucose regulation model is derived with supervised training based on at least one model of glucose regulation selected from the group consisting of: a differential equation model of glucose regulation and a glucose model comprising a set of coupled equations. In some embodiments, the at least one differential equation model of glucose regulation includes a food source function. In some embodiments, the food source function comprises at least one neural network trained with training data selected from the group of: glycemic responses of a population to predetermined foods, and glycemic responses calculated from data for predetermined foods. In some embodiments, the system further comprises an electronic device configured to generate the food source data. In some embodiments, the system further comprises a memory coupled to the model parameter input and configured to store the personalized parameters.

[0017] In another example, the present disclosure provides a method that can comprise: training, on a plurality of attributes, a first neural network (NN) to impute a first subset of the plurality of attributes from a second subset of the plurality of attributes; training a second NN to predict a target value from the first subset of the attributes and the second subset of attributes; receiving a subset of input attributes of a plurality of input attributes from a subject; using the first NN to impute remaining input attributes in the plurality of input attributes; and processing the first subset of inputs attribute and the remaining input attributes with the second NN to

predict a target value. In some embodiments, the first NN comprises an autoencoder. In some embodiments, the second NN comprises a bidirectional recurrent NN. In some embodiments, the recurrent NN is a long short-term memory NN. In some embodiments, the second subset of the plurality of attributes and the subset of input attributes comprise nutrition data for food; and the predicted target value is a glycemic value.

[0018] In another example, the present disclosure provides a system that can comprise: a subject data input configured to receive input attributes from a subject; a first neural network (NN) trained to impute related attributes from input attributes by randomly selecting attributes from sets of attributes having associated target values, and configured to sequentially receive the input attributes; a second NN trained to predict a target value from related attributes and the input attributes, and configured to receive the input attributes and the related attributes generated by the first NN; and a subject data output configured to output and update a predicted target value from the second NN in response to the application of each input attribute to the first and second NNs. In some embodiments, the system further comprises: a data store configured to store training input attributes and corresponding training target values for training the first and second NNs. In some embodiments, the first NN comprises an autoencoder. In some embodiments, the second NN comprises a bidirectional recurrent NN. In some embodiments, the recurrent NN is a long short-term memory NN.

[0019] In another example, the present disclosure provides a method that can comprise: receiving sensor data from at least one sensor that generates biophysical readings for a subject; by operation of a first neural network (NN), embedding the sensor data to generate embedded values; by operation of a second NN, generating imputed embedded values in response to the embedded values, the imputed embedded values including imputed values corresponding to one or more regions of the sensor data; and normalizing the embedded imputed values to generate imputed values. In some embodiments, the regions do not include data that is usable. In some embodiments, receiving sensor data includes receiving data from a first sensor and a second sensor different from the first sensor; and embedding the sensor data includes concatenating data from the first and second sensors. In some embodiments, the first sensor is a glucose monitor. In some embodiments, the second sensor is a heart rate monitor. In some embodiments, the second NN comprises an autoencoder.

[0020] In another example, the present disclosure provides a system that can comprise: at least one biophysical sensor that generates sensor data having missing regions where biophysical readings are determined to be invalid or missing; a first neural network (NN) configured to embed data values from the at least one sensor to generate embedded values; a second NN

configured to generate imputed embedded values in response to the embedded values, the imputed embedded values including imputed values corresponding to the missing regions of the sensor data; and a normalizing system configured to normalize the embedded imputed values to generate imputed values. In some embodiments, the at least one sensor includes a first sensor and a second sensor different than the first sensor; and the first NN is configured to embed sensor data from the first and second sensors in a same time period into single values. In some embodiments, the first sensor is a glucose sensor. In some embodiments, the second sensor is a heart rate monitor. In some embodiments, the second NN comprises an autoencoder.

[0021] In another example, the present disclosure provides a method that can comprise: receiving a validated data set and a query data set, each data set including data values with labels; by operation of a neural network (NN), classifying the validated data set and query data sets with a probabilistic classifier conditioned on the data set values and target labels; and generating a quality score based on a classification result for all data values of one data set. In some embodiments, the method further comprises generating the query data set, including taking biometric sensor readings with corresponding actions as labels. In some embodiments, the biometric sensor comprises a glucose meter. In some embodiments, the labels comprise food log data. In some embodiments, a distribution of the data values has the form p(X, Y, Z) where X is the input distribution, Y is a categorical target of the probabilistic classifier, and Z varies according to which data set the values belong to. In some embodiments, a classification of the probabilistic classifier takes the form $h(x) = p(z = 1|x, Y=1)$, and z equals 0 if x is from the data set with validated labels and Z equals 1 if x is from the query data set.

[0022] In another example, the present disclosure provides a system that can comprise: a data storage system configured to store data sets including data values with labels, the data sets including at least a validated data set and a query data set; and an electronic system in communication with the data storage system that includes at least one neural network configured as a probabilistic classifier configured to classifying the validated data set and query data sets with conditioned on the data set values and target labels, and a quality section configured to examine a classification value for all data values in the query or validated data set and generate a quality value in response thereto. In some embodiments, the system further comprises: at least one biometric sensor configured to generate data values for the query data set. In some embodiments, the biometric sensor comprises a glucose meter. In some embodiments, the validated and query data sets include blood glucose levels with food logs as labels. In some embodiments, a distribution of the data values has the form p(X, Y, Z) where X is the input distribution, Y is a categorical target of the probabilistic classifier, and Z varies according to

which data set the values belong to. In some embodiments, a classification of the probabilistic classifier takes the form $h(x) = p(X=1|x, Y=1)$, and Z equals 0 or 1 depending upon whether x is from the validated data set or query data set.

[0023] In another example, the present disclosure provides a method that can comprise: storing biophysical sensor signals and logged behavior corresponding to the biophysical sensor signals in a data storage device, the stored data comprising training data; training a neural network on the training data to classify biophysical sensor signals as resulting in target behaviors; receiving input biophysical sensor data; and processing the input biophysical sensor data using the neural network to classify a target behavior that results from the input biophysical sensor data. In some embodiments, the biophysical sensor signals include glucose sensor signals and the logged behavior includes logged food data. In some embodiments, the biophysical sensor signals include heart rate monitor signals and the logged behavior includes logged food data. In some embodiments, the target behavior is predicted food consumption. In some embodiments, the method comprises: acquiring the input biophysical sensor signals with at least one sensor for a subject; transmitting the input biophysical sensor signals to the neural network; and transmitting the target behavior to a device of the subject.

[0024] In another example, the present disclosure provides a system that can comprise: a storage system configured to store training data comprising training sensor data and corresponding behavior data; at least one biophysical sensor configured to generate and transmit subject sensor data; and a behavior prediction system configured to receive the subject sensor data and comprising at least one electronic system comprising a neural network trained as a classifier that classifies the subject sensor data into a target behavior, the classifier trained with the training data. In some embodiments, the training data comprising training sensor data from a plurality of different biophysical sensors; and the at least one biophysical sensor includes the plurality of different biophysical sensors. In some embodiments, the training sensor data includes glucose levels and the behavior data includes logged food corresponding to the glucose level; the at least one biophysical sensor includes a glucose meter; and the target behavior is predicted food ingestion. In some embodiments, the training sensor data includes heart rate data and the behavior data includes logged food corresponding to the heart rate data; the at least one biophysical sensor includes a heart rate monitor; and the target behavior is predicted food ingestion. In some embodiments, the system further comprises: the behavior prediction system is further configured to transmit the target behavior; and a subject device configured to receive the target behavior.

[0025] In another example, the present disclosure provides a method that can comprise: receiving and storing string data corresponding to a description of a food item; applying the string data to a language processor configured to determine nominative words and non-nominative words from the string data; in response to the nominative words, querying an item database with the nominative words; in response to non-nominative words, querying the item database with the non-nominative words; and generating a list of query results in response to the querying, the list of query results comprising recipes for the food items. In some embodiments, the method further comprises: the language processor is further configured to determine nominative words as explicit ingredients; and filtering the responses to the querying with the explicit ingredients to generate the list of query results.

[0026] In another example, the present disclosure provides a system that can comprise: a storage device configured to store a database comprising descriptions of food items; a language processing system comprising at least one computing device configured to process text strings to determine nominative and non-nominative words; a query system comprising at least one computing device configured to apply first queries to the database in response to the nominative words generated by the language processing system and to apply second queries to the database in response to the non-nominative words to generate a list of query results in response to the queries, the list of query results comprising recipes for the food items. In some embodiments, the language processing system is further configured to determine nominative words as explicit ingredients; and the query system is further configured to filter responses to the first or second queries.

[0027] In another example, the present disclosure provides a method that can comprise: receiving and storing first data comprising properties of an item; receiving and storing second data comprising constituents of the item ranked in order of prevalence in the item; determining the properties for at least one of the ranked constituents in a database to generate look-up data; determining at least one amount of the at least one constituent in the item in response to the look-up data; and storing the at least one amount of the at least one constituent as output data. In some embodiments, receiving and storing first data includes receiving nutrition information for a food item; and receiving and storing second data includes receiving ranked ingredient data for the food item. In some embodiments, receiving and storing first and second data includes capturing and processing image data of a food label of the food item. In some embodiments, the first data includes n properties; second data includes m constituents; determining the properties for each constituent includes creating and storing an nxm matrix of constituents and their properties; and determining the amount of each constituent in the item includes solving a system of equations

corresponding to y = Ax, where y is the amount of an ingredient, x is a constituent and A is the matrix. In some embodiments, determining the amount of each constituent includes applying the matrix A to a neural network configured for linear regression analysis.

[0028] In another example, the present disclosure provides a system that can comprise: a data pre-processing section coupled to receive first data comprising properties of an item and second data comprising constituents of the item having a ranked in order of prevalence in the item, and including a processing device configured to create a data structure that represents properties for each constituent; and an analysis section coupled to receive the data structure and including a processing device configured to determine determining the amount of each constituent in the item. In some embodiments, the system further comprises: an input device configured to capture the first and second data for the item. In some embodiments, the input device comprises an image capture device configured to capture the image of a label for the item. In some embodiments, the first data comprises nutrition information of a food item and the second data comprises ingredients of the food item. In some embodiments, the first data includes n properties; second data includes m constituents; the data structure comprises a topological mapping of constituents and their properties; and the analysis section is configured to solve a system of equations corresponding to y = Ax, where y is the amount of an ingredient, x is a constituent and A is the topological mapping. In some embodiments, the topological map is a matrix and the analysis section comprises a neural network configured for linear regression analysis.

[0029] In another example, the present disclosure provides a method that can comprise: receiving and storing first data comprising properties of an item; receiving and storing second data comprising constituents of the item ranked in order of prevalence in the item; determining the properties for at least one of the ranked constituents in a database to generate look-up data; determining at least one amount of the at least one constituent in the item in response to the look-up data; and storing the at least one amount of each constituent as output data.

[0030] In another example, the present disclosure provides a method that can comprise: training a word embedding system having a weighting matrix with training data comprising string descriptions of items and properties of the items to embed the string descriptions of items into an embedded space weighted according to the properties of the items; and applying an input string description of the item to the trained word embedding system to infer an output word embedding weighted according to the properties of the items. In some embodiments, training the word embedding system includes training with food string descriptions with nutrition

information as the properties. In some embodiments, applying the input string includes applying a string description of a food item.

[0031] In another example, the present disclosure provides a system that can comprise: a storage system configured to store training data that includes string descriptions of items and properties of the items; a word embedding system using a neural network trained with the training data to embed words of the string descriptions into an embedded space with a weighting derived from the properties of an item corresponding to one of the string descriptions; and an input configured to receive an input string and apply it to the word embedding system to generate word embeddings weighted according to the properties. In some embodiments, the training data includes word description of food items and nutrition information for the food items. In some embodiments, the input string includes a description of the food item and the generated word embeddings are weighted according to the nutrition information.

[0032] In another example, the present disclosure provides a method that can comprise: training a word embedding system having a weighting matrix with training data comprising string description of items and properties of the items to embed the word string items into an embedded space weighted according to the properties of the corresponding item; and applying an input string description of the item to the trained word embedding system to infer output word embedding with the property weighting.

[0033] In another example, the present disclosure provides a system that can comprise: a storage system configured to store training data that includes string descriptions of items and properties of the items; a neural network configured as word embedding system trained with the training data to embed words of the strings descriptions into an embedded space with a weighting derived from the properties of the item corresponding to the string description; and an input configured to receive an input string and apply it to the trained word embedding system to generate word embeddings weighted according to the properties.

[0034] In another example, the present disclosure provides a method that can comprise (a) obtaining text-based descriptions of a plurality of food items and, for each of the plurality of food items, (i) nutrition data and a glycemic value or (ii) nutrition data or a glycemic value; (b) generating embeddings of the text-based descriptions of the plurality of food items; (c) inferring, based at least on the embeddings, a glycemic value for each food item in the plurality of food items for which a glycemic value was not obtained and nutrition data for each food item in the plurality of food items for which nutrition data was not obtained; (d) training a supervised machine learning algorithm on the nutrition data and the glycemic values of the plurality of food items to predict a glycemic value of a given food item from nutrition data of the given food item.

In some embodiments, the method comprises providing the glycemic value of the given food item to the supervised machine learning algorithm to predict the glycemic value of the given food item. In some embodiments, the glycemic value is a glycemic index or a glycemic load. In some embodiments, (b) comprises applying an unsupervised learning algorithm to the text-based descriptions of the plurality of food items. In some embodiments, the unsupervised learning algorithm is a dimensionality reduction algorithm. In some embodiments, the unsupervised learning algorithm is an n-gram or bag-of-words model. In some embodiments, the supervised machine learning algorithm is a deep neural network.

[0035] In another example, the present disclosure provides a system that can comprise: a data storage system configured to store at least a first database and a second database, the first database including descriptions of first items with corresponding attributes, the second database including descriptions of second items with corresponding target values, at least some of the first items being different than the second items; an embedding system comprising at least a first computing device configured to merge the first and second databases to generate training data that includes merged item descriptions with corresponding attributes and target values; and at least a first inference system comprising a machine learning system trained with the training data to infer target values from attributes. In some embodiments, the descriptions of items comprise word descriptions. In some embodiments, the items are food items, the attributes are nutrition data of the food items, and the target values are glycemic response values. In some embodiments, the glycemic response values are selected from the group of: a glycemic index and a glycemic load. In some embodiments, the system can further comprise a data capture section configured to acquire nutrition data with at least a subject device, and wherein the at least first inference system is configured to infer a glycemic index value from at least the acquired nutrition data. In some embodiments, the system can further comprise at least a second inference system that is configured to determine a blood glucose value of a subject in response to at least glycemic response values of foods indicated as ingested by the subject. In some embodiments, the embedding system comprises at least one neural network configured to embed descriptions of first and second items into an embedded space.

[0036] In another example, the present disclosure provides a system that can comprise: a data acquisition system configured to acquire attribute values for items; and at least a first inference system configured to infer target values from the acquired attribute values, the first inference system including: at least one neural network trained with training data generated by embedding at least a first data set and second data set, the first data set including descriptions of

items with corresponding attributes, the second data set including descriptions of items with corresponding

target values. In some embodiments, the system further comprises a training agent configured to train the at least one neural network with the training data. In some embodiments, the system further comprises at least a second inference system configured to infer a response for a subject from at least inferred target values. In some embodiments, the target values are glycemic response values for food items, and wherein the attribute values are nutrition values of the food items. In some embodiments, at least the attribute values are text values embedded into a vector space. In some embodiments, the system further comprises an application server configured to transmit data to an application executed on a subject device in response to at least the inferred target values.

[0037] In another example, the present disclosure provides a method that can comprise: training a neural network with time series training data of a first modality and time series training data of a second modality to create a first model that generates time series data of the second modality from time series data of the first modality; training a second model with the generated time series of the second modality, time series training data of a third modality, and time series data of a fourth modality to generate time series data of the fourth modality; until a convergence condition is reached, iteratively testing the second model on the time series data of the first modality and the time series data of the third modality; and responsive to reaching the convergence condition, predicting second modality data by testing the second model with data of the first modality. In some embodiments, the method comprises: acquiring the time series training data of the first modality with a first type sensor; and acquiring the time series training data of the second modality with a second type sensor. In some embodiments, the second type sensor is a glucose meter, and wherein the time series data of the second modality includes glucose levels over time. In some embodiments, training the neural network to create the first model includes training with N sets of time series training data, and wherein training the first model with the estimated time series training data of the first modality and time series training data of at least the third modality includes training with M sets of time series data. In some embodiments, the method comprises testing the first model with the N sets of time series data and the M sets of time series data and updating the first model in response to error values of the testing, and wherein the trained first model is the first model with the smallest error. In some embodiments, reaching a convergence condition includes calculating an error value not greater than a threshold.

[0038] In another example, the present disclosure provides a system that can comprise: an initial model section that includes a first model trained to generate time series data of a second modality from time series data of a first modality with M sets of training data; a training section

that includes: a second model derived from the first model and configured to generate time series data of at least a third modality from at least time series data of a fourth modality with N sets of training data, and a testing section configured to test the second model with the M and N sets of training data, and update the second model in response to test error values; and an inference model that is the second model with the lowest test error value, configured to infer time series data of the second modality from time series data of the first modality. In some embodiments, the first model, the second model and the inference model comprise neural networks. In some embodiments, the time series data of the first and second modalities are biophysical sensor data. In some embodiments, at least the time series data of the first and second modalities are glucose levels corresponding to glucose meters. In some embodiments, the third and fourth modalities are glucose levels. In some embodiments, the training section comprises: an inverse model that is an inverse of the first model and configured to generate estimated time series data of the first modality from the time series data of the third and a fourth modality; an estimator section configured to generate linear parameters from the estimated time series data of the first modality and the time series data of the third modality; section configured to generate mapped time series data of the first modality from time series data of the third modality using the linear parameters, wherein the second model is trained with the mapped time series data of the first modality.

[0039] In another example, the present disclosure provides a method that can comprise: training a neural network with time series training data of a first modality and time series training data of a second modality to create a first model that generates time series data of the second modality from time series data of the first modality; until a convergence condition is reached: using a second model to generate estimated time series data of the first modality from a mixture of time series data from a third modality and a fourth modality, wherein the second model is initiated as an inverse model of the first model; using the estimated time series data of the first modality and time series data of the third modality, training the second model to estimate linear fitting parameters; using the estimated linear fitting parameters to generate analogous time series data of the first modality from the time series data of the third modality; linearly mapping the analogous time series data of the first modality to the time series data of the third modality; training a third model using the linearly mapped analogous time series data from the first modality mixture of time series data of the third modality and time series data of the fourth modality to generate a mixture of time series data from the third modality and time series data from the fourth modality, wherein the third model is an inverse of the second model; modifying the second model to be an inverse model of the third model; and evaluating whether

the convergence condition has been reached. In some embodiments, training the third model includes initializing the third model as the first model.

[0040] In another example, the present disclosure provides a method for training a neural network to calibrate time series data. The method can comprise receiving calibrated time series data for a biophysical response and corresponding raw time series data for the biophysical response; training, on the calibrated time series data and the corresponding raw time series data for the biophysical response, a neural network to generate calibrated time series data, which training comprises updating parameters of the neural network based on a difference between (i) an output of the neural network for a given raw time series and (ii) a corresponding calibrated times series; receiving raw input time series data generated by a biophysical sensor; and generating calibrated time series data by applying the raw input time series data to the neural network. In some embodiments, the raw input time series data is generated by a glucose meter. In some embodiments, the neural network is trained to cancel drift present in the raw input time series data. In some embodiments, the raw time series data and raw input time series data are generated by glucose meters. In some embodiments, training the neural network further comprises domain specific feature engineering. In some embodiments, training the neural network comprises unsupervised training.

[0041] In another example, the present disclosure provides a method that can comprise: building data structures from a plurality of data sets having an ordering, the data structures including interval trees based on the ordering; determining if any structures have missing intervals in the interval tree; if a data structure has a missing interval, creating data for the missing interval by imputing data values for the missing interval; accessing the data structures by at least searching the interval trees in response to query data; and forming a tabular data structure from the accessed data values that includes a column reflecting the ordering. In some embodiments, the data sets comprise actions ordered in time. In some embodiments, determining if any of the data structures have missing intervals includes classifying data structures into a first class if they have no missing intervals and a second class if they have missing intervals. In some embodiments, accessing data values from the data structures includes an operation selected from the group consisting of: selecting a data structure for a query operation; querying a region of a data structure dictated by the ordering; joining query results; and merging overlapping regions of different data structures. In some embodiments, forming the tabular data structure includes forming a dataframe from the accessed data values. In some embodiments, forming the tabular data structure includes forming a dataframe from the accessed data values. In some embodiments, the data sets comprise different subject events having an ordering, and wherein

forming the tabular data structure includes forming a tabular data structure that includes different subject events over a queried time period. In some embodiments, at least one of the subject events is a biophysical response of the subject. In some embodiments, the biophysical response is a glucose level of the subject.

[0042] In another example, the present disclosure provides a system that can comprise: a data store configured to store tabular data sets, each having data values with an ordering; and memory comprising machine-executable instructions that when executed by a processor cause the processor to perform operations comprising: create data structures that include interval trees based on the ordering, determining if any of the interval trees includes missing intervals, if an interval tree has a missing interval, imputing data for the missing interval, accessing data values from the data structures by at least searching the interval trees of the data structures in response to query data, and forming a tabular data structure from the accessed data values that includes a column reflecting the ordering. In some embodiments, the data store is configured to store tabular data sets having time or date column corresponding to subject actions. In some embodiments, the processing section is configured to execute an operation selected from the group consisting of: selecting a data structure query operation; querying a region of a data structure dictated by the ordering; joining query results; and merging overlapping regions of structures. In some embodiments, the data store is configured to store tabular data sets comprising different subject events having an ordering, and wherein the processing section is configured to forming tabular data structures that includes different subject events over a queried time period. In some embodiments, at least one of the subject events is a biophysical response of the subject. In some embodiments, the biophysical response is a glucose level of the subject.

[0043] Another example of the present disclosure provides a non-transitory computer readable medium comprising machine executable code that, upon execution by one or more computer processors, implements any of the methods above or elsewhere herein.

[0044] Another example of the present disclosure provides a system comprising one or more computer processors and computer memory coupled thereto. The computer memory comprises machine executable code that, upon execution by the one or more computer processors, implements any of the methods above or elsewhere herein.

[0045] Additional aspects and advantages of the present disclosure will become readily apparent to those skilled in this art from the following detailed description, wherein only illustrative embodiments of the present disclosure are shown and described. As will be realized, the present disclosure is capable of other and different embodiments, and its several details are capable of modifications in various obvious respects, all without departing from the disclosure.

CLAIMS

WHAT IS CLAIMED IS:

1. A computer-implemented method for training a reinforcement learning algorithm, the method comprising:

(a) providing a reaction model comprising a personalized learned body model configured to generate a simulated biophysical or behavioral response of a subject responsive to a plurality of inputs;

(b) until a convergence condition is achieved, iteratively:

(i) generating, using the reinforcement learning algorithm, a recommendation comprising a recommended meal or physical activity,

(ii) processing the recommendation using the reaction model to generate a predicted response of the subject to following the recommendation, wherein the predicted response comprises a biophysical or behavioral response,

(iii) applying a first reward function to the predicted response to generate a first reward value indicative of a benefit or detriment of the predicted response toward a health measurement of interest; and

(iv) training the reinforcement learning algorithm in a first stage using the first reward value,

(c) providing the recommendation to the subject;

(d) measuring the health measurement of interest of the subject responsive to following the recommendation;

(e) applying a second reward function to the measured health measurement of interest to generate a second reward value; and

(f) training the reinforcement learning algorithm in a second stage using the second reward value.

2. The method of claim 1, further comprising encoding the measured health measurement of interest of the subject into a low-dimension latent space for providing to the second reward function.

3. The method of claim 1, wherein the first reward function is the same as the second reward function.

4. The method of claim 1, wherein generating the predicted response of the subject further comprises:

processing encoded historical data of the subject using a trained predictor configured to infer the predicted response;

processing the recommendation using an adherence model configured to evaluate an adherence of the subject to the recommendation; and

selectively processing outputs of the trained predictor and the adherence model using the personalized learned body model.

5. The method of claim 6, wherein generating the predicted response further comprises processing the simulated biophysical response using an autoencoder and a generative adversarial network.

6. The method of claim 1, wherein the convergence condition comprises a criterion based at least in part on the magnitude of the first reward.

7. The method of claim 1, wherein the reinforcement learning algorithm comprises a neural network.

8. The method of claim 1, wherein the first reward function generates the first reward at a higher rate than the second reward function generates the second reward.

9. The method of claim 1, wherein the reaction model is trained on historical biophysical or behavioral response data from the subject.

10. The method of claim 9, wherein the reaction model is updated during inference based on new biophysical or behavioral response data of the subject.

11. The method of claim 1, wherein the reaction model comprises a supervised machine learning model.

12. The method of claim 1, wherein the first reward function and the second reward function generate a positive reward when the recommendation results in a beneficial predicted response or measured health measurement of interest of the subject.

13. The method of claim 1, wherein the first reward function and the second reward function generate a negative reward when the recommendation results in a detrimental predicted response or measured health measurement of interest of the subject.

14. A system comprising one or more computer processors and one or more storage devices having instructions stored thereon that are operable, when executed by the one or more computers, to cause the one or more computer processors to perform operations comprising:

(a) providing a reaction model comprising a personalized learned body model configured to predict a biophysical or behavioral response of a subject responsive to a plurality of inputs;

(b) until a convergence condition is achieved, iteratively:

(i) generating, using the reinforcement learning algorithm, a recommendation comprising a recommended meal or physical activity,

(ii) processing the recommendation using the reaction model to generate a predicted response of the subject to following the recommendation, wherein the predicted response comprises a biophysical or behavioral response,

(iii) applying a first reward function to the predicted response to generate a first reward value indicative of a benefit or detriment of the predicted response toward a health measurement of interest; and

(iv) training the reinforcement learning algorithm in a first stage using the first reward value;

(c) providing the recommendation to the subject;

(d) measuring the health measurement of interest of the subject responsive to following the recommendation;

(e) applying a second reward function to the measured health measurement of interest to generate a second reward value; and

(f) training the reinforcement learning algorithm in a second stage using the second reward value.

15. A non-transitory computer storage medium having instructions stored thereon that are operable, when executed by one or more computer processors, to cause the one or more computer processors to perform operations comprising:

(a) providing a reaction model comprising a personalized learned body model configured to predict a biophysical or behavioral response of a subject responsive to a plurality of inputs;

(b) until a convergence condition is achieved, iteratively:

(i) generating, using the reinforcement learning algorithm, a recommendation comprising a recommended meal or physical activity,

(ii) processing the recommendation using the reaction model to generate a predicted response of the subject to following the recommendation, wherein the predicted response comprises a biophysical or behavioral response,

(iii) applying a first reward function to the predicted response to generate a first reward value indicative of a benefit or detriment of the predicted response toward a health measurement of interest; and

(iv) training the reinforcement learning algorithm in a first stage using the first reward value;

(c) providing the recommendation to the subject;

(d) measuring the health measurement of interest of the subject responsive to following the recommendation;

(e) applying a second reward function to the measured health measurement of interest to generate a second reward value; and

(f) training the reinforcement learning algorithm in a second stage using the second reward value.

16. The method of claim 1, wherein the predicted response comprises the biophysical response.

17. The method of claim 16, wherein the predicted response comprises a predicted glucose response, and wherein the health measurement of interest comprises a glucose level.

18. The method of claim 17, further comprising measuring the glucose level of the subject using a continuous glucose monitor.

19. The method of claim 1, wherein the predicted response comprises the behavioral response.

20. The method of claim 1, wherein training the reinforcement learning algorithm in the first stage or the second stage comprises adjusting a set of weights or parameters of the reinforcement learning algorithm.

21. The method of claim 1, further comprising using the personalized learned body model to generate the simulated biophysical or behavioral response of the subject in real-time.

22. The method of claim 21, wherein the simulated biophysical or behavioral response of the subject is generated based at least in part on real-time sensor data of the subject.

23. The method of claim 21, further comprising:

determining if the simulated biophysical or behavioral response of the subject is outside of a pre-determined range; and

if the simulated biophysical or behavioral response of the subject is outside of the pre-determined range, providing the recommendation to the subject in real-time, wherein the recommendation is selected to adjust an actual biophysical or behavioral response of the subject to be within the pre-determined range.