

MACHINE LEARNING BASED PREDICTIVE SYSTEM FOR ENHANCING CONSUMER AND BUSINESS DECISIONS USING DIETARY HABITS OF INDIVIDUALS

BACKGROUND

Field of the Invention:

This invention relates to the field of learning user dietary habits through machine learning algorithms and providing efficient decision options for both, the consumer and the business entities based on the predictions from the collected user data.

Description of the Related Art:

- Food and beverage (F&B) sector business establishments are readily adopting digital technologies and means to increase their sales and customer feedback. It is not uncommon in a modern society to find restaurants offering their menus on a tablet computer instead of paper menus. Moreover, the tablet computers allow for an automated process of food selection by the consumer, leading to lesser waiting times and easy integrated process of order placement and payment.
- However, such a system lacks the intelligence to guess a consumer's dietary preferences even if they are regular in their dining habits. Therefore, it is important to develop intelligent solutions that not only help automate the process but also provide a more personalized experience by learning and adapting to the dietary preferences of an individual. By learning the dietary preferences over time through registering the eating habits and/or other factors, an intelligent system can be developed that can help consumers in the process of restaurant selection based on their location and preferences.
- The problem with such a system is the multiplicity of food items and their names across restaurants and geographical regions around the world. Some existing solutions try to extract and learn the semantics behind each restaurant's menu and food listings by the presence of words like "chicken", "fish", etc. However, this process is not exhaustive and prone to errors in learning as well as predictions due to the uncertain nature of the data structure in the first place.
- It is, therefore, necessary to develop a unified backend dictionary of machine learning identities to represent any food item across any geographical region and name variations. This is one of the problems that is tackled through the methods and systems provided in the current invention.
- Other existing systems for learning user habits for food recommendations are based on counting the occurrences over time. However, such a system will also fail with increasing number of business establishments in the central database and the multiplicity of food items.

SUMMARY

- The present invention provides methods and system design for developing algorithms that can learn certain important parameters from the dietary habits of consumers. These parameters further contribute to the process of inferring certain metrics for the business as well as the consumers.
- In any embodiment of the present invention, three subsystems are primarily defined: consumer, vendor, and restaurants. All the consumers reside in the consumer subsystem and have access to certain exposed functionalities of the system defined in this patent. Vendors are business owners that have multiple food (dining) establishments within their operational domains and form the vendor and restaurant subsystems respectively. The invention pertains to methods and systems designed and exposed to each of the defined subsystems and the resulting analytics from them used to enhance the productivity of this entire ecosystem.
- In any embodiment of the present invention, the consumers are provided with a downloadable application on their mobile devices. The application requires the consumers to answer certain general questions about their age, gender, ethnicity, lifestyle choices, and/or others. These parameters are static inputs to the algorithms running on the cloud servers. Specifically, these parameters are utilized to make an assumption for a prior distribution of food preferences over a dataset of food items.
- In any embodiment of the present invention, the consumers use the said application on their mobile devices to select and order food in restaurants within our business agreements. The advantage of such a process is twofold: first, the consumers get access to an automated system that enables easy food selection, and second, the business entities can track various parameters related to the volume of consumers that are active within their setup at any point in time. As a result of this process optimization, intelligent analytics are generated that help a business establishment (necessarily in dining domain) adapt to and in all possibility, optimize their functionality for improving profits.
- In one embodiment of the present invention, the business entities have the provision for a certain intelligent queue management system. This system allows the consumers to pre-order food options through their mobile applications. The queue management system tracks the location of consumers and notifies the business entities to initiate food preparation process based on the consumer's location or indicated time of arrival to eat. The methods and systems of the current invention then compute and estimate each consumer's waiting/stall time between online ordering and the actual arrival to the restaurant.
- In any embodiment of the present invention, the system stores the user habits in form of a mixture distribution of certain kind. This mixture distribution has at least four components, each pertaining to breakfast, lunch, pre-dinner, and dinner. The system tracks the user habits over a 24 hours period for occurrences and patterns of their ordering history. This form of learning and representation has two advantages: first, the database throughput is largely decreased as the consumers' preferences are learned and represented through simple predefined functional notations hence simplifying the underlying data structure, and second, the method is semi-supervised leading to an efficient estimate of each consumer's habits and food preferences.
- In any embodiment of the present invention, the system incrementally updates the parameters of the said mixture distribution. This update may result in generation of new mixture components (within defined maximum upper bounds), extinction of certain inactive components according to a decay model or a total replacement of current

mixture with a uniform distribution, implying that any particular item is preferred at any point of time in the day by a specific consumer.

- In any embodiment of the present invention, the learned distributions of consumers can be used for time-based prediction of the amount of expected traffic in a certain business establishment, the demand for popular items, and/or others thus leading to an efficient process of inventory and queue management by the business entities.
- This Summary was provided to introduce a selection of concepts in a simplified form that is further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

- FIG. 1 shows the ecosystem within which the application framework is operable, namely, vendors, consumers and restaurants/food stalls subsystems.
- FIG. 2A shows the database structure in which machine learning dataset is stored and queried using its hash key.
- FIG. 2B depicts an alternate database structure model for storing the said machine learning dataset, having two keys for faster indexing, storage and query operations.
- FIG. 3 is a diagrammatic description of the basic data structure (DS) of the e-menu across vendors.
- FIG. 4 describes the valid and invalid relationship models for the e-menu data structure.
- FIG. 5A is a diagrammatic description of the method in which each consumer's dietary habits are tracked using a mixture model.
- FIG. 5B shows an example of how the learned mixture model can be used to infer a consumer's dietary preference at any point of time within a 24-hour day by business entities.

- FIG. 6 is a diagrammatic description showing the static and dynamic input parameters to the machine learning algorithm that are used by the cloud machine learning APIs to generate consumer and business related predictions and analytics.
- FIG. 7 shows the data structure of a consumer as stored in the database.
- FIG. 8 is a flowchart depicting the process flow starting from the consumer login in the mobile application, then placing an order, followed by updating the machine learning parameters and ending with the confirmation of item delivery.
- FIG. 9 shows the application interface as presented to a consumer allowing selections and inputs in a user-friendly manner.

DETAILED DESCRIPTION

[001] For the purposes of providing an understanding of the principles of the invention and the scope of applications involved therein, reference will now be made to the embodiments illustrated in the drawings and specific language will be used to describe the same. It will nevertheless be understood that no limitation of the scope is thereby intended. Any alterations and further modifications in the described embodiments, and any further applications of the principles as described herein are contemplated as would normally occur to one skilled in the art.

[002] In the following description, algorithms and functions may be shown in block diagram form in order not to obscure the present invention in unnecessary detail. Conversely, specific application implementations shown and described are exemplary only and should not be construed as the only way to implement the present invention unless specified otherwise herein. Additionally, block definitions and partitioning of logic between various blocks is exemplary of a specific implementation. It will be readily apparent to one of ordinary skill in the art that the present invention may be practiced by numerous other partitioning solutions. For the most part, details concerning timing considerations and the like have been omitted where such details are not necessary to obtain a complete understanding of the present invention and are within the abilities of persons of ordinary skill in the relevant art.

- [003] The invention describes methods and system design for learning users' dietary habits over time using a machine-learning algorithm that requires no supervision, minimal database storage and access to these variables in an optimal manner through the data structures developed therein. The invention enables consumers and business entities to improve their decision-making process in a manner that increases the consumer experience for an end consumer and optimizes the functionality of a business establishment leading to improved profit margins.
- [004] FIG. 1 shows the ecosystem within which the application is operational. This ecosystem comprises of three subsystems namely, vendor subsystem **102**, stall/restaurant subsystem **108**, and the consumer/user subsystem **104**. From this point onwards, the terms restaurant, stalls, and canteens will be understood to refer to the same entity, that is the subsystem **108**. Likewise, the terms users and consumers will be understood to refer to subsystem **104**.
- [005] The vendor subsystem **102** comprises of business owners depicted by **106A**, **106B**, and **106C**. The definition of vendors encompasses business owners that have a certain number of restaurants within a geographical area. Various elements of the current invention provide said vendors with overall analytics and relevant business predictions about their owned business establishments. For instance, as shown in the figure, vendor **106B** has restaurants **110A**, **110B**, and **110C** within its subsystem. A certain subset of activities in the stall subsystem **108** is required to be logged into the database of the vendor subsystem also.
- [006] In many embodiments of the current invention, the vendors and stalls have access to said functionalities provided by the cloud APIs **100** through a web link with authentication.
- [007] In other embodiments of the current invention, the vendors and stalls are provided access through a stand-alone application installed on their mobile or other computing devices.
- [008] The user subsystem **104** comprises consumers with access to a downloadable application on their mobile devices. This subsystem is exposed to a certain set of actionable functions that act as input to the machine learning algorithms in the cloud servers. The user subsystem provides static data inputs in form of parameters including but not limited to their age, gender, ethnicity, and/or others, and dynamic data inputs related to their food habits. The term food habit is understood to refer to the particular dietary preferences of any user, for instance, vegetarian/vegan/non-vegetarian, times at

which users are active on the application to order food items, the frequency of ordering certain food items across restaurants, and/or others.

[009] The ecosystem comprises cloud APIs **100**, that enable database storage/retrieval and maintaining the machine learning algorithm parameters for each consumer. The cloud APIs are accessible through communication links **112**. In most embodiments, the links are operable through frameworks like REST and/or protocols like MQTT, XMPP, Web Sockets, and/or others.

[0010] In the preferred embodiment of the current invention, MQTT channels are defined for establishing a communication link between the consumers and vendors/stalls. The use of MQTT allows for a low power, low latency messaging solution as compared to long polling, HTTP or XMPP methods. The communications are necessarily performed with a quality of service (QoS) level of 1, implying that the messages are delivered at least once within a certain maximum time period. As a result, each MQTT message updates a particular table in the database allowing users to check for and restore previous session activity.

[0011] In a typical embodiment of the invention, users can place their order request in a particular stall within the stall subsystem. This action is simultaneously updated by the cloud APIs in the database of the respective vendor as well as the particular restaurant. The operation of updating the database along with providing notifications to the concerned restaurant of intended activity allows restoration of user sessions if the internet connection is disrupted due to any reason in this process. For instance, consider a failure scenario where the user requested particular food items from a stall and the e-connection was disrupted with the user not receiving any request acknowledgments. In this case, upon re-connection, the framework will perform a check for latest session status and bring the current session into that particular state automatically.

[0012] In typical implementations, the vendor subsystem has primarily view only rights to the business related data collected over their particular stall subsystem. This implies that vendors can access data to perform certain analytics on collected data without requiring any input to the machine-learning model. Likewise, the stalls in stall subsystem have partial write access facilitated through the communication links. This access is provided in form of actionable buttons or links on their user interfaces allowing them to interact with consumers in the user subsystem. For instance, the restaurants will have to indicate through an action in their user interface, if a particular e-order by an end customer is accepted or rejected, estimated time to completion, and/or others. Inputs from the stall

subsystem are used primarily to indicate completion of transactions, which in turn can trigger the machine-learning algorithm to update the database. In other implementations, the machine-learning algorithm may be triggered by an actionable input from the user subsystem's user interface. This form of algorithm trigger is termed as passive triggering approach and saves computations.

[0013] In another embodiment of the current invention, the machine-learning algorithm may be actively triggered at regular time intervals to allow dynamic push update notifications to a consumer if needed. An active learning scheme as such will result in increased database access, however, may be preferred for increasing business transactions through means of notifying active users of deals that the system calculates may result in a certain amount of profit if activated.

[0014] FIG. 2 depicts two possible paradigms in which the machine learning input parameters are defined. The current invention assumes that for each restaurant in the stall subsystem, the food menu or the e-menu is prepared against a fixed list of pre-decided categories and tags. This kind of storage and inferring system backend simplifies the architecture by obviating the need to perform semantic analysis of menus across various restaurants. The pre-processing of data is necessary and helpful since a particular cuisine may have multiple variations of the same food item referred by various names as decided by the restaurant. To simplify the meaning of menu items across restaurants, a dictionary of food categories and their local names is created and updated constantly. At any point in time, this dictionary must contain references to any food item listed at any stall within the ecosystem.

[0015] In any embodiment of the current invention, the database being used by the machine learning APIs can be a relational database (RDBMS) or a non-relational database (noSQL). In the preferred embodiments, a noSQL database backend is used to simplify data storage in form of a javascript object notation (JSON) data structure. Equivalently, having a JSON formatted data structure allows the use of common JSON parsing libraries (lodash, underscore, and/or others) to find patterns in data and subsequently update the machine learning parameters.

[0016] As shown in FIG. 2A, each food item is uniquely represented by an identifier referred to as the *ml_id* **120**. This identifier is used by the hashing function to query the table and access other records **126A** and **126B** like its associated category *cat* **122** and the name of food item *name* **124A** and **124B**. Each record **126A** and **126B** is then represented by its primary key *ml_id* and associated parameters. By storing the data in

this format, where *ml_id* is the hash key, each food item can be individually queried for its properties.

[0017] However, in a rapidly increasing number of dictionary items in the machine learning key database, the data may alternatively be managed in form as depicted in FIG. 2B. In this implementation, the food item's category *cat* is the hash query key **128** along with another sub-query key **130** that is the *ml_id*. Records are then uniquely represented by the combination of their category and identifiers. This kind of database approach helps retrieve items within a certain category with only a single query operation.

[0018] In another embodiment, the dictionary of machine learning keys can be stored in form of an XML file, JSON file, and/or other data formats allowing easy retrieval and parsing of the records therein.

[0019] In most embodiments of the current invention, an API is provided for an administrator to create e-menus of various business entities within the ecosystem. The API provides functionalities that allow the creation of a new item listing in the dictionary if not already present, refer to an existing food category *cat* with the creation of new *ml_ids* if not already present, and referring to an already existing food item if a record is already present in entirety. Checks may be performed at regular intervals to detect any discrepancies within the central key database and e-menus across the business entities, though the possibility of such an error happening is decidedly very less as a result of the said API being provided to the system administrators.

[0020] Turning now to FIG. 3, the data structure for storing the e-menu **140** of restaurants is shown. This e-menu is necessarily created by a system administrator instead of requiring the restaurants to fill in the information. The e-menu key points to various categories of food listings in the restaurant, **142A**, **142B** and **142C**. Each of these categories represents a broad encapsulating term, for instance, "mains", "desserts", "drinks", and/or others. This categorization is primarily for user accessibility enhancement. Each such category comprises an array of JSON objects **144**, each representing a particular food listing of the restaurant. Each such object comprises fields including but not limited to its identity in the machine learning dictionary *ml_id*, the display name of food listing as preferred by the restaurant *name*, and the price in units dependent on the region of application deployment.

[0021] In the preferred embodiment, this data structure is stored as a JSON format value indexed by the unique identity of the particular restaurant. The e-menu data structure is accessible by the applications running on the mobile devices of consumers in the user

subsystem. The application is responsible for parsing the obtained e-menu data structure and displaying it in a manner that allows selection.

[0022] In other implementations, the array objects **144** may also incorporate an object *options*, that represents various states in which the particular food listing can be requested in. For instance, the *options* object can refer to the amount of sugar level in a drink, selection of particular spices in a food item and so on.

[0023] Storing the e-menu in this format helps in identification of intersecting patterns between any user's learned dietary habits or preferences and the e-menu data structure across restaurants in the stall subsystem **108**. In embodiments where the e-menu is stored in other formats, appropriate conversion functions need to be invoked by the machine learning algorithm to trace common interest areas. In other embodiments, the restaurants are given write access to the *name* and *price* fields of the defined e-menu data structure.

[0024] Turning now to FIG. 4, relationships between the *ml_id* and *cat* fields of the dictionary is provided. Each *ml_id* can be linked to only one category of food items. For instance, *ml_id1* **150** can have a link **162** with category *cat1* **156**. However, a simultaneous link **160** with category *cat2* **158** is not possible and is a void or invalid link. Hence *ml_id* to *cat* links are one to one functions, however, *cat* to *ml_id* links can be many to one function. For instance, **158** can have valid links **164** and **166** from **152** and **154**.

[0025] Using the above-mentioned relationship model in the data structure definition allows each menu item across various restaurants to have a unique reference (unless the item has the same name as well as category). At the same time, such a model captures the variations of a dish across cuisines and e-menus of different restaurants in the stall subsystem.

[0026] FIG. 5A and FIG. 5B describe the primary method used in the machine learning backend of cloud servers. The horizontal axis **178A** and **178B** are times of the day represented with a resolution of minutes from 1 to 24 hours. The vertical axis **170A** and **170B** represents all the item categories in the dictionary defined in FIG. 2. Each consumer's habit in the user subsystem can be represented by such a 2-dimensional graph of distributions across time and item categories.

[0027] When a user orders a particular food item in any restaurant, the distribution is initialized with one component having the mean as that point of time and a certain variance decided from the prior data provided by the consumer. The reason for choosing a Gaussian function to represent the action of placing an order is that the consumer can

order the same item on different days with a certain distribution across time. Since user habits are usually clustered around their breakfast, lunch, supper and dinner times, it is assumed that each consumer can exhibit at most four components in the mixture across time and item categories as represented by **176**. Therefore, for any item with a *cat* value (belonging to category) of cat_k , the probability of it being ordered at any point of time in the day is represented as $\sum_{i=1}^4 \alpha_i^k G^k(\mu_i, \sigma_i)$, where $\sum \alpha_i = 1$, and $G(\cdot)$ is a Gaussian. Usually, a food item will not be ordered constantly across the four time divisions per day as depicted by **174** and hence, we assume the distribution to have an upper bound of four components. However, if at any time, the number of components increases beyond four, the item is said to have a uniform distribution across the time, implying no preference or dependency on time as represented by **172**. Such items are generally usual items like water, tea or coffee.

[0028] It is noted that the category of the food item is used by the algorithm for learning the distributions and not the *ml_ids*. This allows for a more generalized learning process instead of a very specific model where learning may not be possible due to the multiplicity of food items across various restaurants. In the preferred embodiment, a list of such specific food item *ml_ids* can be stored within the category's description to indicate preferred items within the learned categories.

[0029] In the preferred embodiment, the frequency of ordering a certain category of items is used to incrementally update the mixture distribution parameters. This implies that if a particular item category is constantly ordered, its mixing proportion will be altered to reflect the preference. Various methods could be used to determine if a particular item has a distribution over time periods of each day or it is a general purpose item having a uniform distribution. In one implementation, entropy of the distribution may be used to estimate its closeness to a uniform distribution. In another implementation, mutual information, cross entropy or the KL divergence of the calculated distribution with a uniform distribution can be used to determine a decision to represent the item as uniformly distributed.

[0030] In another embodiment, the mixture distributions are associated with a decay model implying that if the particular item was not ordered frequently across days, the component's mixing proportion decreases. This implies that a particular food item category was ordered as a one-time choice, and is not usually preferred by the consumer

on a day-to-day basis, thus leading to very less information gain if the machine-learning algorithm represented this choice in its functions.

[0031] In the preferred embodiment, the machine-learning model is updated passively each time the consumer orders something. However, such an approach has disadvantages related to the fact that at any point of time, the system is one-step (the last time consumer placed an order) behind the current user scenario. This is inconsequential if no decay model is applied. In other embodiments, the machine learning system is updated at regular intervals for all the users in the database. Such an active update scheme will cause increased server loads and may be activated for certain users only based on certain criteria.

[0032] It is noted from FIG. 5A that there is a discontinuity in representation across various days, that is, if a food item were ordered at 23:59 hours, the distribution will not span next day. In order to eliminate this discontinuity, the distribution functions may be circularly shifted from the 1 to 24 hours timescale depicted.

[0033] In the preferred embodiment, each user's non-zero components are stored and the rest are not integrated into the database. By storing only the non-zero components, the minimum amount of space is utilized by the user subsystem resulting in increased efficiency. A typical data structure describing the machine learning model for any user then assumes the form {Distribution type, [values]}, where distribution types define if it is mixture or uniform, followed by the parameter values if it is a mixture distribution across time.

[0034] FIG. 5B depicts a Bayesian inference scenario where inference can be made about any user's preferred item choice with time. This question is represented as $f(item_i | t \in [t_1, t_2])$ where $f()$ represents the value of preference of certain item within the time period as depicted by the shaded region **180**. A simple *max* operation across items on the vertical axis can then provide user preference within the time range specified. This inference can then be further used to notify the user of availability of the item within the restaurant and/or other notifications or actions.

[0035] In one embodiment, this learned data structure is also stored locally on a consumer's mobile device. In this scenario, all calculations for updating the model are performed on the user's device actively or passively, followed by updating the cloud database with the computed values.

- [0036]** In preferred embodiments, the cloud calculates certain priors over food categories depending on user inputs like their age, gender, ethnicity, religions and/or others. Calculation of a prior implies storing a prior belief of how probable a certain user is about ordering a certain food item. The prior belief functions can be estimated from user surveys or manual assignment. For instance, persons with vegetarian preference will never order non-vegetarian food category options and any such singular activity if happening in future; can be discarded as a one-time event with no importance towards the machine-learning model for the particular user.
- [0037]** In yet another implementation, the system can use time-normalized values of each user to determine the popularity of items across various age groups, ethnicity, gender, and/or others. These statistics can, in turn, be utilized to enhance the prior distributions of any user and enhance the learning strategies of the machine-learning model.
- [0038]** By using the learning model as described, it is possible to infer quantities such as an individual's fitness, total calorie intake, and/or others. For the vendors, analytics like the geographical popularity of their restaurants, preferred e-menu items across their restaurants, expected monetary gains, and/or others can be generated.
- [0039]** In some embodiments, an intelligent queue management system (QMS) is also integrated that triggers certain alarms based on the proximity of any consumer from the restaurant. Such a system allows the business establishments to reorder their preparation queue in terms of individuals who are more likely to reach the destination earlier. This kind of approach is contrasted against a purely first come first serve basis, which may not always be logical for a cloud-based mobile application of sorts as described in this invention.
- [0040]** In some embodiments, consumer feedback about service may also be obtained to improve the machine learning algorithm parameters. For instance, consumers may indicate that a particular food recommendation generated by the algorithm is not correct. Depending on this binary feedback, the algorithm may change its parameters to reflect these changes.
- [0041]** Turning now to FIG. 6, the broad input parameters **190** supplied to the machine-learning algorithm on cloud servers **100** and the output parameters are shown. The algorithms generate analytics for consumers **208** and business entities **218**. As described earlier, the machine-learning algorithm can trigger these predictions actively or passively.
- [0042]** The input parameters to the algorithm can be divided into static priors **198** and dynamic inputs **192**. Static prior parameters are quantities such as a consumer's age **226**, gender

200, ethnicity **202**, religion **204**, lifestyle **206**, and/or others. These parameters are important for the machine-learning algorithm as they indicate certain statistical biases about dietary habits of the individuals. For instance, people of a particular geographical region will prefer food items belonging to that area with a high possibility. Likewise, people belonging to a certain religion will constrain themselves to certain food categories. People following a particular lifestyle, for example, fitness oriented, will not consume food items with calorific values higher than a certain threshold. All these assumptions are purely statistical and not definitive and hence are best represented through a distribution that reflects this uncertainty and variation across individuals. These prior assumptions and definitions are generally fixed across all consumers in the preferred embodiment.

[0043] In another embodiment, the priors may be specific to each consumer. In this implementation, consumer feedback might be used as binary decision latent inputs that cause a particular prior to alter its distribution across the dictionary of food items.

[0044] In another embodiment, the definition of priors for each variable may be stored and updated on the mobile user's devices instead of calculating on the cloud.

[0045] The dynamic parameters **192** are consumer choices made with every order. These parameters include the food item choice **194** and the time of day when the particular item was selected **196**. Unlike static parameters that are responsible for updating the prior distributions, dynamic parameters update the mixture distributions for any consumer. Selection of food item of a particular category increases the machine learning algorithm's belief of the person liking any items belonging to that category. Time of selection is important as it quantifies whether the particular item is consumed during breakfast, lunch, pre-dinner or dinner. This kind of semi-supervised learning method precludes the need of manually dividing a consumer's eating times. As described earlier, normalization across the time axis (**178**) or item axis (**170**) results in the estimation of different quantities.

[0046] The machine learning algorithm generates consumer predictions **208**, that may include estimated quantities as calorie intake per day/month **210**, health monitoring by way of tracking eating times and calorific intake **212**, dietary suggestions **214** based on static and dynamic parameters chosen, restaurant suggestions **216** based on user's dynamic parameters or current geographical location, and/or others.

[0047] The machine learning algorithm generates predictions **218** relevant to the business entities such as their expected consumer traffic at any time of day/month **220**, popular

item categories across parameters such as age groups of people, ethnicity of people, gender, and/or others. As a result of these statistics, the business entities may make a more informed decision about their inventory and resource management **224**.

[0048] Turning now to FIG. 7, the consumer data structure **240** is shown. This data structure is incorporated in many embodiments to facilitate data access, storage and update by the machine-learning algorithm. The database comprises various users in the ecosystem **242**, where each user's static prior parameters are stored, followed by their dynamic parameters. The dynamic parameters are an array of JSON format structures of format as described earlier: {Distribution type, [values]} **244**. Each record is accessible by a unique *user_id* for every individual in the database **240**.

[0049] FIG. 8 shows the process flow of how the machine learning algorithm is integrated within the food selection and ordering system. To facilitate learning, data must be organized and presented in a manner that is intelligible to the business owners, system admins (of cloud services), as well as the consumers. All these three levels of the hierarchy have different abstractions for the manner of data presentation. Consumer activity is central to the functioning of the machine learning algorithms as they are the primary input data source.

[0050] The process starts at **250** by providing an interface to the consumer for login and authentication. Following this, the user may be required to *e-enter* or select a particular restaurant of choice. At this stage, the system can also provide restaurant recommendations based on the learned machine-learning model and the geographical location of the consumer.

[0051] After the above step, food recommendations are populated at step **252** by utilizing the learned user dietary habits and the categories of food items present in the e-menu of the restaurant. This step is primarily intended to affect a consumer's decision in favor of food items that they like with a high probability. Moreover, these steps provide a more personalized user experience by way of knowing or understanding the time based dietary preferences of each consumer. In some embodiments, the system may also keep track of the success/failure rate of the presentation of each recommendation to the consumer. This latent feedback can be used by the system to further refine future suggestions.

[0052] At step **254**, the user confirms a particular list of selection from the e-menu. This can be a combination of the recommended food items or newer items. As a result of this step, the cloud APIs update the machine learned model parameters for that consumer if

a passively triggered embodiment is employed. In the case of actively triggered embodiment, the user's order history data structure is updated, which the system can use to modify the learned parameters on the next active trigger time instance. Following this step, the appropriate restaurant might be notified of consumer's intent, followed by an approval/denial response. Any further payment triggers may be activated post this response by the business entities.

[0053] On an affirmative response, the mobile application is subscribed to actively receive notifications about the status of order **258**. It is noted that the user can use the mobile application to pre-place the order for his intention of reaching a particular business entity within an estimated stated time. In some embodiments, the machine learning algorithms on the cloud can also track the delays of every consumer between their intent indication and actual arrival time. This data can be used to further notify the business entities about the expected arrival time if location based data is not available or implemented.

[0054] The process flow ends at step **260** when a consumer has reached the intended destination. In some embodiments, this step may additionally trigger a feedback from the user about the quality of performance by the application. The response may be logged for use by the machine learning APIs to alter their learned parameters and provide better recommendations.

[0055] Turning now to FIG. 9, the mobile application provided to the consumers is shown. Step **270** is accessible post user authentication through means in **250**. Step **270** is one of the embodiments where vendor suggestions depending on consumer's location are not generated; instead, a list of business entities is shown in any order. In cases where there is no vendor entity, the application skips directly to step **272**.

[0056] Step **272** shows various restaurants within the ecosystem of vendor selected in step **270**. At this stage, depending on the machine-learned parameters, suggestions may be provided to consumers on which particular listed entities provide a certain food item of their liking.

[0057] Step **274** involves selection of food items from the e-menu listing and recommendations generated by the cloud APIs. Following this step, a consumer may finalize the transaction at step **276** that is further processed by the cloud algorithms for active notifications and learning.

WHAT IS CLAIMED

1. A machine learning scheme implemented on the cloud providing methods and system for:

Representation of user's dietary habits across time using a distribution function and their mixtures in some proportions.

Updating the distribution functions and their mixing proportions passively based on user actions of placing an order for food items in business establishments dealing in food and beverage business.

Storing the learned dietary preference for each user in functional object data format instead of individual preference values for each item.

Deciding whether a particular item has some preference over time or is a general purpose commodity that can be ordered at any time of the day.

In the above case being true, representing the food item's preference distribution as a uniform distribution across the 24-hour day with a discretization of minutes.

Requiring certain consumer inputs like age, gender, ethnicity, age, lifestyle, and/or others and deciding a prior distribution based on such answers. The prior distribution is further used to refine the process of learning and prediction.

2. An implementation of claim 1 where the time axis for representation is discretized to the level of hours.
3. An implementation of claim 1 where the time axis is discretized to the level of seconds.
4. A method of claim 1 where the said distribution functions have an associated decay function that causes the preference of any item to decay with time if it were not ordered at certain closely spaced time intervals.
5. A method of claim 4 where the decay function is modeled in terms of a Poisson distribution for counting occurrences.
6. An implementation of claim 1 where the system is partially or completely implemented and values stored on a consumer's mobile device instead of the cloud servers.
7. An implementation of claim 1 where the machine learning algorithm's parameters are updated actively at regular time intervals based on order history of individuals.

8. A method of claim 1 that allows sending push notifications to consumers about the availability of their choice of food items in a certain business establishment.
9. A computer implemented method providing methods and system for calculation of or inferring auxiliary variables such as:

The fitness level of consumers based on the functions generated as a result of claim 1.

The preferences of a particular food item/category across the different age range of consumers within a geographical region.

The preferences of a particular food item/category across consumers with different ethnicity, gender, and lifestyle choices.

Estimating the expected profit/sales turnover per month based on the data collected for each consumer by claim 1.

Estimating the peak times of sales for each business establishment through the data collected by claim 1.

10. A method of claim 9 providing means for intelligent inventory management from the machine learning data of consumers active within the restaurant in certain durations of time.
11. A mobile device application providing a user interface to facilitate methods and systems defined by claim 1 and 9, and having the provision to:

Select a restaurant of choice from an ordered list and *e-enter* the restaurant for deciding food choices.

Select a restaurant of choice from a list created according to the current geographical location of the consumer.

Select food items of choice from an e-menu listing provided through a combination of claim 1 and the business establishment.

Request approval and indicate expected time of arrival to the restaurant after food items selection.

Subscribe to push notifications indicating the status of an approved order.

12. An intelligent queue management system that provides for:

Maintenance of online queues through alarms based on a consumer's proximity to the restaurant.

Estimation of a consumer's speed of approach towards the restaurant based on time between crossings of certain radii from the restaurant's latitude and longitudinal position.

An intelligent system that tracks and stores the stall times of consumers for each business establishment from their order history for forming a prior belief of their expected time of arrival.

13. A computer implemented method providing a system the means to:

Create an e-menu for business establishments by comparing the display names of food items with their categories.

Maintain a dictionary of machine learning unique ids of food items indicating their category and specific type.

Append to the said dictionary if a new food item is found to have an entirely new category.

Update an item in the said dictionary with a new id if it already belongs to one of the existing categories.

Performing periodic checks in the database of said dictionary for machine learning to check for discrepancies and missing key elements.

14. A computer implemented database for storing the machine learning parameters of claim 1 in a JSON or XML format data structure.

15. A method of claim 1 where consumers can choose to turn off features for tracking and storing their dietary habits.

16. A graphical interface accessible to business establishments as a standalone application or a web interface, allowing:

Read only access to auxiliary metrics derived from data obtained by claim 1, the metrics being quantities such as the volume of consumers per day/month, or popular food items.

Machine learning methods for regression analysis of estimated cash flow per month or day.

Analytical methods to incorporate new food items in their e-menu that may lead to a certain calculated profit based on food preferences of consumers active within the geographical region.