

(19) **United States**(12) **Patent Application Publication**
Kil(10) **Pub. No.: US 2008/0147440 A1**(43) **Pub. Date: Jun. 19, 2008**(54) **EVIDENCE-BASED MEDICINE
SUPERCHARGER****Publication Classification**(51) **Int. Cl.**
G06Q 50/00

(2006.01)

(52) **U.S. Cl. 705/2**(57) **ABSTRACT**


Apparatuses, computer media, and methods for supporting health needs of a consumer by processing input data. An integrated health management platform supports the management of healthcare by obtaining multi-dimensional input data for a consumer, determining a health-trajectory predictor from the multi-dimensional input data, identifying a target of opportunity for the consumer in accordance with the health-trajectory predictor, and offering the target of opportunity for the consumer. An outcome study for a medical treatment from a medical publication is detected. The medical treatment is mapped to a diagnostic and procedural code and a database for health data is accessed using the diagnostic and procedural code. An outcome metric for the medical treatment with a consumer group is associated, and a utility function is generated from the plurality of outcome metrics, where the utility function gauges an efficacy of at least one intervention channel for a consumer.

(75) **Inventor: David H. Kil, Santa Clara, CA (US)**

Correspondence Address:

BANNER & WITCOFF, LTD.**ATTORNEYS FOR CLIENT NO. 005222****10 S. WACKER DRIVE, 30TH FLOOR****CHICAGO, IL 60606**(73) **Assignee: Accenture Global Services
GMBH, Schaffhausen (CH)**(21) **Appl. No.: 11/672,423**(22) **Filed: Feb. 7, 2007****Related U.S. Application Data**

(63) Continuation of application No. 11/612,763, filed on Dec. 19, 2006.






FIG. 1




FIG. 1A




FIG. 2




FIG. 3

400





FIG. 4






FIG. 6

600




FIG. 7

EVIDENCE-BASED MEDICINE SUPERCHARGER

[0001] This application is a continuation of common-owned, co-pending U.S. application Ser. No. 11/612,763 ("Integrated Health Management Platform") filed on Dec. 19, 2006 naming David H. Kil, the entire disclosure of which is hereby incorporated by reference.

FIELD OF THE INVENTION

[0002] This invention relates generally to healthcare management. More particularly, the invention provides apparatuses, computer media, and methods for supporting health needs of a consumer by processing input data.

BACKGROUND OF THE INVENTION

[0003] The U.S. healthcare industry is a \$2T economy with the rate of growth far exceeding that of general inflation. With the aging global population, the current healthcare crisis is expected to worsen, threatening the health of global economy. The existing healthcare ecosystem is zero-sum. The recent pay-for-performance (P4P) experiment by the National Health Services in the United Kingdom resulted in mixed outcomes with incentive-based payments far exceeding the budget with uncertain improvements in patient health. On the other hand, a recent study on the sophistication of healthcare consumers reveals that there is little correlation between consumers' perception of care and the actual quality of healthcare delivered as measured by RAND's 236 quality indicators. Furthermore, given the high churn rate and the propensity of employers to seek the lowest-cost health plan, payers are motivated to focus primarily on reducing short-term cost and carving out the cream-of-the-crop population, resulting in perverse benefit design.

[0004] In healthcare, predictive models are used to improve underwriting accuracies and to identify at-risk members for clinical programs, such as various condition-centric disease management programs. Unfortunately, predictive models typically use year-1 payer claims data to predict year-2 cost. Some predictive modeling vendors predict future inpatient or emergency-room episodes since they represent high-cost events. The emphasis on cost makes sense given that the impetus for predictive models came from private and government payers struggling with rising healthcare costs.

[0005] Evidence-based medicine (EBM) is an attempt to apply scientific evidence to making care decisions for patients. A lot of EBM guidelines are derived from medical journals, where teams of researchers rely on randomized controlled trials and observational studies to draw inferences on the efficacy of various treatments on carefully selected patient populations. Pharmacovigilance or study of adverse drug reactions is an example of EBM.

[0006] Current EBM vendors, such as Active Health Management, a wholly owned subsidiary of Aetna, and Resolution Health, rely on a team of physicians reading and codifying relevant medical journals. The resulting EBM database is applied to population claims data consisting of medical, Rx, and lab claims data in order to identify patients not receiving proper EBM guidelines, i.e., with "gaps" in treatment. Physicians of the identified patients are contacted through faxes or telephone calls with instructions or recommendations on how to close the gaps in treatment. A number of shortcomings

exist with the current EBM implementation. Many EBM studies suffer from small sample size, thus making generalization difficult and sometimes inaccurate. A corollary of the first shortcoming is that most EBM studies are at a selected population level and do not provide drilldown information at a sub-population level. That is, if not everyone benefits from an EBM guideline, it may be dangerous to apply the guideline to the entire study population, which begs for a careful tradeoff between specificity and sensitivity. Guidelines typically do a poor job of translating study outcomes into metrics that end stakeholders care about. For example, payers pay a particular attention to cost, which is not the same as improving surrogate endpoints that are therapeutic in nature with various time frames for healing or outcomes improvement. Publication bias and conflicting results encourage ad hoc decision making on the part of payers in the area of utilization management, such as coverage denials and medical necessity reviews. Furthermore, relying on published guidelines discourages the use of autonomous or loosely guided search for anomalies or precursors to adverse outcomes using a large of amount of integrated data assets and intelligent search algorithms based on machine learning.

[0007] Clearly, there is a desperate need for an integrated solution for providing healthcare management.

BRIEF SUMMARY OF THE INVENTION

[0008] The present invention provides apparatuses, computer media, and methods for supporting health needs of a consumer by processing input data.

[0009] With one aspect of the invention, an integrated health management platform supports the management of healthcare by obtaining multi-dimensional input data for a consumer, determining a health-trajectory predictor from the multi-dimensional input data, identifying a target of opportunity for the consumer in accordance with the health-trajectory predictor, and offering the target of opportunity for the consumer. Multi-dimensional input data may include claim data, consumer behavior marketing data, self-reported data, and biometric data.

[0010] With another aspect of the invention a consumer is assigned to a cluster or clusters based on the multi-dimensional input data. A characteristic of the consumer may be inferred from a subset of the multi-dimensional input data.

[0011] With another aspect of the invention, a cluster is associated with a disease progression, where the cluster is associated with at least one attribute of a consumer. A target of opportunity is determined from the cluster and the disease progression. An impact of the target of opportunity may be assessed by delivering treatment to a consumer at an appropriate time.

[0012] With another aspect of the invention, a target of opportunity is extracted from medical information using a set of rules for the multi-dimensional input data.

[0013] With another aspect of the invention, a previous event that occurred before a subsequent transition event is identified. A correlation between the previous event and the subsequent transition event is measured from historical data to assign multidimensional strength or utility indicators to a discovered rule.

[0014] With another aspect of the invention, an enrollment healthcare selection for the consumer is recommended based on multi-dimensional input data.

[0015] With another aspect of the invention, an outcome study for a medical treatment from a medical publication is

detected. The medical treatment is mapped to a diagnostic and procedural code and a database for health data is accessed using the diagnostic and procedural code. An outcome metric for the medical treatment with a consumer group is associated, and a utility function is generated from the plurality of outcome metrics, where the utility function gauges an efficacy of at least one intervention channel for a consumer.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] The present invention is illustrated by way of example and not limited in the accompanying figures in which like reference numerals indicate similar elements and in which:

[0017] FIG. 1 shows an architecture of an integrated health management (IHM) platform in accordance with an embodiment of the invention.

[0018] FIG. 1A shows a Service Oriented Architecture (SOA) framework in accordance with an embodiment of the invention.

[0019] FIG. 2 shows a method of determining multimode health-trajectory predictors in accordance with an embodiment of the invention.

[0020] FIG. 3 shows a flow diagram for an evidence-based medicine supercharger in accordance with an embodiment of the invention.

[0021] FIG. 4 shows a flowchart for an autonomous healthcare data exploration system in accordance with an embodiment of the invention.

[0022] FIG. 5 shows an illustrative conceptual example of the optimal health benefit design in accordance with an embodiment of the invention.

[0023] FIG. 6 shows an example of Markov modeling of assessing a target of opportunity in accordance with an embodiment of the invention.

[0024] FIG. 7 shows computer system 100 that supports an embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

Integrated Health Management Platform in a Service Oriented Architecture (SOA) Framework

[0025] FIG. 1 shows an architecture of integrated health management (IHM) platform 100 in accordance with an embodiment of the invention. IHM platform 100 creates for payers a virtuous circle of integrated informatics leading to improved and real-time decision making leading to healthier members leading to improved profitability and cost savings leading to improved market share. For consumers who must share an increasing burden of medical costs, the execution of IHM platform may lead to improved health and subsequent cost savings. In the following discussion, a consumer may be an employee of a company or an individual choosing a healthcare plan from a plurality of plans and consuming products/services from various healthcare stakeholders. An objective of the consumer is to maximize benefits associated with good health by choosing a healthcare plan that is "best" for the individual or his/her family and improving health through timely preventive and proactive health actions.

[0026] The IHM platform consists of the following four components:

[0027] 1. Multimode health-trajectory predictors module 101: Instead of focusing on predicting future cost alone using claims data as most predictive models do now, multimode health-trajectory predictors leverage claims data

153, self-reported data 151, and consumer behavior marketing data 155, coupled with inference engines 115, to provide a comprehensive set of future attributes useful to assess the level of impact through various consumer-engagement channels. Claims data 153 may include medical claims, pharmacy claims, prior authorization, and lab results (e.g., blood tests) for a consumer. Consumer-engagement channels may encompass secure e-mails, Interactive Voice Recording (IVR) calls, cellphone text messages, and nurse calls. Data Merge & Cleaning 109 performs extract-transform-load (ETL) of disparate data assets to form a consumer-centric view while cleaning data prior to weak-signal transformation through digital signal processing (DSP) and feature extraction 111. Disease clustering and progression module 113 subsequently forms disease clusters and estimates disease progression probabilities. Clustering optimization & inference 115 performs clustering using attributes that are meaningful from the perspective of predicting future health trajectories and impactability with the inference engine filling in unobserved variables using the instantiated variables. A modular predictive model is developed for each consumer cluster so that a collection of locally optimized predictive models can provide a globally optimal performance 117. Finally, a set of health scores encompassing health scores, behavior/lifestyle scores, engagement scores, impact scores, data-conflict scores, cost scores, and clinical scores is output 119.

[0028] 2. Targets-of-opportunity finder 103: Leveraging consumer-understanding technologies, an evidence-based-medicine (EBM) supercharger (shown as EBM supercharger 300 in FIG. 3), and an autonomous insight crawler, one can identify targets of opportunities in various consumer touch points. The four major opportunities lie in clinical gaps 121, treatment adherence 123, lifestyle/behavior 125, and psychosocial parameters 126. Impact assessment is made based on the aggregate future impact of all the identified targets of opportunities 127.

[0029] 3. Resource-allocation manager 105: Resource-allocation manager (RAM) 105 funnels the right members to the right consumer touch points at the right time by maximizing multi-objective functions. Also included in RAM 105 are consumer-understanding technologies and iterative benefit design borrowing salient concepts from adaptive conjoint analysis, predictive modeling, and Pareto multi-objective optimization. Furthermore, mixing-in currently available technologies into consumer touch points in conjunction with dynamic progressive content tailoring allows one to go beyond the typical nurse-based care model, which is inherently not scalable especially with the projected worsening nurse shortage in the labor market. (Resource-allocation manager 105 is Pareto efficient if no consumer can be made better off without another consumer being made worse off.) The fundamental idea here is building a multi-objective constrained optimization engine 137 as a function of consumer, intervention-channel, benefit-program profiles (129, 131, 135) and utility functions 133 derived from the impact analysis engine.

[0030] 4. Impact-analysis engine 107: This module tells one what works for which population segments, by how much, and why in a drilldown mode. It facilitates the use of utility functions in the framework of resource-allocation optimization as done in defense battlefield resource management. The methodology employed uses predictive

modeling, combinatorial and stochastic feature optimization with respect to outcomes, and propensity-score shaping. After selecting candidate population for analysis **141**, one performs thorough matching in the two-dimensional space of propensity and predictive scores **143** to create control and intervention groups **145** for an “apple-to-apple comparison.” One then create rules of engagement for statistically significant outcomes, which are further validated through focus-group study **149** and survey using the minimum number of necessary questions **148**. Validated rules **150** are stored in the master rules database for production implementation.

[0031] The four above components **101-107** complement one another and are ideally suited to assessing the incremental benefits of bringing new data assets and business processes into enterprise operations. In order to facilitate integration into and compatibility with typical payer enterprise applications, the IHM implementation (e.g., IMH platform **100**) adheres to an enhanced Service Oriented Architecture (SOA) framework. A key idea here is maximizing synergy among business process primitives, data models, and algorithm models so that one can reduce latency between the generation of actionable knowledge and its production implementation.

[0032] FIG. 1A shows Service Oriented Architecture (SOA) framework **160** of IHM platform **100** in accordance with an embodiment of the invention. Framework **160** increases synergy in data models, mathematical models, and business-process models that are important in ensuring the success of IHM Platform **100**. Inputs **161** consist of data library **163**, algorithm library **165**, and business-process libraries **167**, which get updated with the latest discoveries. The processing layer uses the building blocks of business processes and algorithms tailored to underlying data models to produce intermediate processing outputs as well as actionable insights that feed to multimedia outputs for dissemination to the key stakeholders.

[0033] Data library **163** includes Consumer Touch Points (CTP) **169**, Utilization Management (UM) **171**, Underwriting Questionnaire (UWQ) **173**, D&B: Dun & Bradstreet (D&B) database **175**, Electronic Medical Records (EMR) **177**, and Health Risk Assessment (HRA) database **179**.

Multimode Health-Trajectory Predictors

[0034] FIG. 2 shows process **200** for determining multimode health-trajectory predictors in accordance with an embodiment of the invention. Instead of focusing on cost prediction, multimode health-trajectory predictors attempt to understand current and predict transitions in Bayesian relationships among the many semi-orthogonal outcomes attributes so that one can maximize positive impact through delivering the right intervention touch points to the right consumers at the right time before adverse transitions occur.

[0035] In healthcare, predictive models are used to improve underwriting accuracies and to identify at-risk members for clinical programs, such as various condition-centric disease management programs. Typically, prior art predictive models predict year-2 cost using year-1 payer claims data. Some prior art predictive modeling vendors predict future inpatient or emergency-room episodes since they represent high-cost events. The emphasis on cost makes sense given that the impetus for predictive models came from private and government payers struggling with rising healthcare costs.

[0036] Focusing on cost alone ignores the complex, multifaceted nature of healthcare consumers. Knowing future cost with R-sq of 10-25% is different from being able to impact the future health trajectory of each consumer. For example, it may be more beneficial to touch John suffering from pre-diabetic conditions with body mass index (BMI) of 32 than to intervene on behalf of Mark who has to go through kidney dialysis three times a week because of end-stage renal disease. From a cost perspective, Mark may be 20-40 times more expensive. But from an impact perspective, John would be a better candidate because his current conditions are more amenable to actions that can be taken now to prevent unpleasant consequences in the near future.

[0037] As a result of cost emphasis, prior art predictive models extract a standard set of features from Rx and/or medical claims data and apply linear regression or neural networks to predict year-2 cost. Typical features include disease flags and year-1 summary cost and utilization statistics, such as average inpatient cost per month, average Rx cost per month, # of physician visits per month, etc. Some predictive models divide the population into sub-groups using inputs from clinicians with the goal of designing a model tailored to each sub-group (MedAI). However, it may be quite difficult to design optimal clusters given the complexities of and interplays among the many factors that determine future health trajectories.

[0038] In order to address the shortcomings of the current generation of predictive models, an embodiment of the invention incorporates the following concepts:

[0039] 1. Use of input data such as claims data **251**, self-reported data **255**, consumer behavior marketing (CBM) data **253**, and biometric data **257** is augmented with inference engine **209** to predict multiple semi-orthogonal attributes with the goal of finding the best way to engage and motivate healthcare consumers to create positive impact. Input data is typically provided by electronic health record (EHR) database **203**. Not everyone will have all the data assets. Therefore, key unknown variables need to be estimated using inference engine **209**.

[0040] 2. Flexible-dimension clustering process **211** creates an optimal set of consumer clusters from an impact perspective instead of using the same old disease hierarchy to create disease-centric consumer clusters.

[0041] 3. Adaptive hypermedia content creation **221** leverages a comprehensive understanding of consumer needs and how to best provide a positive impact.

[0042] As shown in FIG. 2, inputs include:

[0043] Claims data **251**: It is comprised of Rx/med/lab data, utilization-management (UM) data including pre-authorization, Rx/med benefit data, program touch-point data, Web log data, and limited member demographic data.

[0044] Consumer behavior marketing (CBM) data **253**: This externally purchasable data provides inferred behavior, lifestyle, and attitudinal information on consumers from their demographic data and credit history.

[0045] Self-reported data **255**: This includes health risk assessment (HRA), ecological momentary assessment (EMA), and experience sampling method (EMA) data administered through multiple communication channels, such as the Internet, cellphone, set top box, etc.

[0046] Biometric data **257**: This encompasses data from wearable sensors (Bodymedia's BodyBugg™, Nike+ shoe sensor, polar band) and attachable sensors (glu-

cometer, blood-pressure cuff, spirometer, etc.) transmitted through wired or wireless networks.

[0047] As shown in FIG. 2, processing includes:

[0048] Mixer **201**: Not everyone will have all the data elements. Therefore, mixer **201** organizes incoming data into a schema appropriate for frame-based dynamical data processing. Furthermore, it differentiates between 0 and an empty set ϕ .

[0049] Preprocessing **205**: This step performs secondary data audit and consumer-centric data structure generation. Primary data audit occurs during data creation in enterprise data warehouse (EDW).

[0050] 1) Data audit: Outliers are normalized using multi-pass peak-shearing. Multiple debit/credit entries and ghost claims are eliminated. It looks for potential gender/age mismatch errors (grandmother or father giving birth to a baby or a premature baby's neonatal claims being assigned to his or her parents) using a look-up table.

[0051] 2) Consumer-centric data structure generation: For each consumer, we create an efficient data structure from memory and processing perspectives. It is a hierarchical structure encompassing the entire consumer touch-point suite of channels.

[0052] Transform **207**: This step creates various band-pass-filtered maps over time. For instance, International Classification of Disease (ICD) 9/10 codes from medical claims and National Drug Codes (NDC) from Rx claims are converted into hierarchical condition-versus-time maps to facilitate the analysis of disease progression and the creation of disease clusters. Moreover, such a representation can help one to infer behavioral patterns from linking discrete events or following medication adherence for managing chronic conditions. A combination of ICD and Current Procedure Terminology (CPT) codes is used to derive Milliman & Robertson (M&R) categories over time, which is useful in assessing the utilization of various service types (inpatient, outpatient, emergency room, physician office visit, etc.) over time. Biometric data is processed through a large number of transformation algorithms, such as the fast Fourier transform, wavelet transform, local cosine transform, ensemble interval histogram, etc., in order to glean locally dynamic behaviors over time. Due to the infrequent nature of HRA and CBM data (i.e., people do not change their behavior or lifestyle every hour), locally dynamic behaviors serve as anchor points that vary much more slowly so that one can investigate the cumulative effects of linked local events over time on behavior change. The entire transformation process is analogous to multi-rate signal processing. At the end of transform, we extract a large number of static and dynamic features from each transformation space, as well as higher-order linked attributes spanning multiple transformation spaces in order to glean insights into disease clustering, disease progression, and their interplay with the consumer's psychosocial behavioral traits.

[0053] Inference engine **209**: Knowing certain unobserved traits can be quite useful in devising tailored intervention strategies. Let x_{claims} , x_{CBM} , x_{SR} , and x_{bio} represent the four data sets as previously discussed. If knowing one's body mass index (BMI) is desirable, one first builds modular predictive models from the sub-population that has BMI data such that $P(BMI|x_{claims})$,

$P(BMI|x_{CBM})$, etc. constitute a feasible set of models for predicting BMI conditioned upon having other data assets. This model can be in the form of Bayesian networks, regression or classification algorithms leveraging parametric and non-parametric learning algorithms.

[0054] Flexible-dimension clustering **211**: This is an iterative process leveraging multiple fitness functions and predictive models as part of clustering. This step generates a set of clusters for each outcomes variable such that the output dispersion compression is maximized for improved prediction accuracy.

[0055] 1) For each outcomes variable, one performs feature optimization to find a sufficient-statistics feature subset.

[0056] 2) One performs clustering using k-means, expectation-maximization (EM), and Kohonen's self-organizing feature map. After clustering, there are N_C clusters for each outcomes variable. For each cluster, one calculates the dispersion σ_i , $i=1, \dots, N_C$ of each of the outcomes distributions and compare it with the overall dispersion σ_T from the entire population. The dispersion-compression ratio (DCR) $r_i = \sigma_T / \sigma_i > \gamma$, where $\gamma > 1$, is a predetermined dispersion-compression threshold for accepting the i^{th} cluster based on its ability to compress the outcomes distribution. One creates a set of samples that pass the DCR test.

[0057] 3) For the samples that do not pass the first DCR test, repeat steps 1-2 until there is no sample left or the number of remaining samples is less than the minimum sample size.

[0058] Automatic model calibration **215**: In real-world problems, data characteristics remain rarely stationary over time. With process **200**, step **213** determines whether training is needed to update process **200** for new medical developments. For example, introduction of new medical technologies and drugs, changes in benefit plans and fee-reimbursement schedules, changing demographics, and even macroeconomic cycles can affect data characteristics. Built-into the automatic model calibration algorithm **215** is a data-mismatch estimator that keeps track of statistical parameterization of key data assets over overlapping time frames and consumer clusters after removing secular trends, e.g., medical-cost inflation. Model parameters are updated and stored in model parameters database **217**. During model initialization and subsequent re-calibration, the following takes place:

[0059] 1) Perform preprocessing step **205**, transform step **207**, inferring step **209**, and flexible dimension clustering step **211**

[0060] 2) Feature optimization for each consumer cluster and outcomes variable using combinatorial and stochastic algorithms

[0061] 3) Model performance tuning to find the point of diminishing returns

[0062] 4) Multiple-model combining

[0063] Multiple-model scoring **219**: Once process **200** has been trained, multiple-model scoring **219** is performed for input data **251-257**. One generates the following health scores:

[0064] 1) Health scores as a function of current chronic conditions and predicted disease progression

[0065] 2) Behavior and lifestyle scores computed heuristically as a function of reported, observed (medica-

tion adherence, frequent ER visits, the level of interaction with care-management nurses, etc.), and inferred behavior and lifestyle attributes

[0066] 3) Engagement scores as a function of reported, observed, and inferred psychosocial and collaborative-filtering parameters

[0067] 4) Impact scores working in concert with evidence-based-medicine (EBM) supercharger 300 and utility functions associated with targets of opportunities and derived from the impact-analysis engine

[0068] 5) Conflict scores as a function of discrepancies between reported and observed behavioral/lifestyle factors and claims data

[0069] 6) Cost scores for multiple future time periods in chronic vs. acute categories

[0070] 7) Clinical utilization scores in terms of inpatient, emergency room/urgent care centers, medication, etc.

[0071] Adaptive hypermedia content generation 221: This module generates a tailored report of 1-2 pages succinctly summarizing current health conditions, likely future states, targets of opportunities, action plan, and benefits with drilldown menu.

[0072] In accordance with an embodiment of the invention, process 200 provides outputs including:

[0073] Consumer-centric metadata for a comprehensive view (both tabular and scientific visualization) of each consumer with appendages linking consumer-centric metadata to various stakeholders to facilitate stakeholder-centric data transformation

[0074] Health scores

[0075] Adaptive hypermedia content tailored to each consumer

Evidence-Based Medicine Supercharger

[0076] FIG. 3 shows a flow diagram for evidence-based medicine (EBM) supercharger 300 in accordance with an embodiment of the invention. From an EBM guideline or a medical journal article 351, evidence-based-medicine supercharger 300 generates a set of multidimensional inferred and observed utility functions, which is an essential ingredient in developing optimal resource allocation strategies. The utility function can be multidimensional at multiple levels of granularity in terms of patient or consumer clusters, leading to an $M \times N$ matrix, where M and N represent the number of utility components or objectives and the number of consumer clusters, respectively. For example, consumer clusters generated from the health-trajectory predictors may encompass the following groups: (1) those who are generally healthy from a claims perspective, but exhibit poor health habits in terms of high BMI and "couch-potato" characteristics; (2) those who suffer from chronic illnesses amenable from a lifestyle intervention, such as diabetes and cardiovascular disease; (3) people who have multiple co-morbid conditions, but one cannot find treatment-related claims records ($N=3$). From a segmented drilldown impact analyses of three intervention channels (Interactive Voice Response (IVR), health behavior coaching, and case management ($M=3$)), one determines that the most effective intervention channels for the three population clusters are (IVR, health behavior coaching), (health behavior coaching), (case management and health behavior coaching), respectively. The utility function is a 3×3 matrix,

where each element x_{ij} contains a utility score or return on investment for the i^{th} intervention channel applied to the j^{th} consumer cluster.

[0077] In accordance with an embodiment of the invention, evidence-based-medicine supercharger 300 includes:

[0078] Input databases:

[0079] 1) EBM database 317: It consists of EBM rules, taxonomy for inducing rule parameters from medical journal, population parameters, rule strength, mapping look-up tables that map condition and drug names to ICD-9 and NDC, respectively, and utility function. Population parameters encompass inclusion and exclusion criteria. Rule strength is a function of publication rank using a page-ranking algorithm, author prestige based on the number of connections in the publication network, journal prestige based on the number of circulation, sample size, percentage of total cost affected, longitudinal duration, and the number of corroborating articles. The EBM taxonomy facilitates efficient induction of EBM-rule parameters from an exemplary journal abstract as shown in the Appendix. More algorithmic details will be discussed in the processing-algorithm subsection.

[0080] 2) Electronic Health Records (EHR) 203: This database contains claims data 251, self-reported data 255, and consumer behavior marketing (CBM) data 253.

[0081] Processing algorithms

[0082] Text mining 301: The Appendix shows a semi-structured abstract from an article published in the New England Journal of Medicine. Instead of using a bag-of-words or natural-language-processing feature vector and a Naïve Bayes classifier to rank an abstract, one simply detects whether an abstract reports an outcomes study or not. This is a much easier problem and defers the strength-of-evidence classification until after integrated outcomes analysis. Next, one uses a combination of key words, $tf*idf$ text weights (in which the importance of a word is based on its frequency of occurrence in a document and normalized by its natural frequency of occurrence in a corpus) with stemming and stop words, and distance measures from key words to fill in the hierarchical tree EBM database fields in the areas of:

[0083] 1) Type of outcomes research

[0084] 2) Patient characteristics: size, dropout rate (if available), characteristics in terms of inclusion and exclusion criteria, longitudinal duration, and trigger criteria

[0085] 3) Reported results

The distance measures are necessary to leverage lexical analysis to understand higher-level relations and concepts between words in a sentence or a paragraph.

[0086] Automatic EBM rule induction 303: Given the EBM database fields extracted from a medical journal, one uses secondary look-up tables to map drug names, diagnoses, and procedures onto NDC, ICD-9, CPT-4, and laboratory codes commonly used in claims-payment systems.

[0087] Human-Computer Interface (HCI) for human confirmation 305: The induced EBM rule along with the highlighted abstract is presented to a clinician for final confirmation with or without edit.

[0088] EBM population identification **307**: One identifies potential control and intervention populations using the inclusion, exclusion, and trigger criteria. The presence or absence of the trigger criteria assigns a patient to the intervention or control group, respectively, provided that the patient satisfies the inclusion and exclusion criteria.

[0089] Dual-space clustering **309**: This step creates meaningful consumer clusters that are homogeneous in the optimized baseline-period-attribute-and-outcomes (y) vector space. The baseline period equals the pre-intervention period of a fixed duration

[0090] 1) For each EBM guideline, one builds models that predict various outcomes metrics. Associated with each predictive model is an optimal feature subset ($X \in \mathbb{R}^N$, where N is the optimal feature dimension) derived from a combination of stochastic and combinatorial optimization algorithms.

[0091] 2) In the vector space spanned by X, one performs clustering using k-means, expectation-maximization (EM), and Kohonen's self-organizing feature map algorithms. After clustering, there are N_c clusters. For each cluster, one calculates the dispersion σ_i , $i=1, \dots, N_c$ of each of the outcomes distributions and compare it with the overall dispersion σ_T from the entire population. The dispersion-compression ratio (DCR) $r_i = \sigma_T / \sigma_i > \gamma$, where $\gamma > 1$, is a predetermined dispersion-compression threshold for accepting the i^{th} cluster based on its ability to compress the outcomes distribution for more precision in applying EBM from an outcomes perspective. One creates a set of accepted samples for which clusters in X are sufficiently precise for performing integrated outcomes analysis. One selects the clustering algorithm that provides the highest DCR.

[0092] 3) For the remaining population samples, perform feature optimization to derive a new optimal feature subset $X^{(k)}$. Compress $X^{(k)}$ into X_c ($\dim(X_c) \ll \dim(X^{(k)})$) using linear discriminant analysis (LDA) and discretized outcomes metrics should they be continuous. Next, perform clustering in the vector space spanned by X_c and y. Prior to clustering, normalize the vector space so that mean and standard deviation of each component will be 0 and 1, respectively. The standard deviation of y can be higher to reflect its importance in determining clusters. Keep the clusters whose DCRs > 1 .

[0093] 4) For the remaining clusters, repeat step iii until the number of remaining samples is below the minimum threshold, i.e., $(k) \rightarrow (k+1)$. The final remaining samples represent the final cluster.

[0094] Integrated outcomes analysis **313**: For each cluster, perform case-controlled impact analysis leveraging predictive and propensity-score models to account for both regression to the mean and selection bias. A comprehensive set of outcomes metrics encompasses both observed and inferred variables. For the inferred variables, we estimate individual and cluster prediction accuracies so that we can assess the level of statistical significance as a function of cluster size and model accuracy.

[0095] Utility function generation **315**: Finally we generate a set of utility functions.

[0096] 1) Two-dimensional marginal utility functions over individual outcomes metrics and population clusters

[0097] 2) One-dimensional utility function over a composite outcomes metric with weights

[0098] 3) Pareto Frontier set for multiple outcomes metrics based on a user-defined multi-objective function

[0099] Outputs of evidence-based-medicine supercharger **300** include:

[0100] Utility functions tailored to each stakeholder, a composite outcomes metric, or multi-objective optimization or Pareto-efficient plots

[0101] Outcomes metrics

Autonomous Healthcare Data Exploration System

[0102] FIG. 4 shows a flowchart for autonomous healthcare data exploration system **400** in accordance with an embodiment of the invention. Autonomous healthcare data exploration system **400** explores healthcare database to look for "interesting" relationships autonomously using various signal processing and data mining algorithms. There is often substantial hidden insight in healthcare data that can be discovered. Autonomous data exploration is sometimes associated with fraud detection. In healthcare, gaming or exploitation of loopholes in fee-reimbursement policies can be a serious problem, which has led to utilization management or medical necessity review by payers. For example, one study reports that 39% of physicians surveyed use at least one of the following three gaming methods:

1. Exaggerating the severity of patients' conditions
2. Changing patients' billing diagnoses
3. Reporting signs or symptoms that patients didn't have

[0103] Fraud detection has been around for over two decades in a myriad of forms. It typically looks for outliers or uses models learned from labeled training data to identify suspicious activities for human confirmation. The two most widely used areas are in credit-card and financial industries. The U.S. Securities and Exchange Commission (SEC) and research boutique firms pore through tick-by-tick financial data to look for anomalous trading patterns that can signal insider trading.

[0104] Just to illustrate the difficulty of transitioning commercial antifraud solutions to healthcare, the U.S. Government Accountability Office reports that instead of adopting commercially available antifraud software to Medicare use, the Health Care Financing Administration (HCFA) chose to enter into a multi-year agreement with the Los Alamos National Laboratory, citing numerous difficulties with adopting commercial software. Unfortunately, no such software—commercial or custom-built—is in widespread use today.

[0105] The focus on fraud pits one stakeholder against another when outright fraud is relatively rare, and a soft form of exploiting system loopholes is more common in healthcare. Therefore, there is a need for a more sophisticated and less demeaning system focused on learning hidden causal relations between treatment and health outcomes (both positive and negative) so as to gain the widest possible acceptance from all the stakeholders.

[0106] FIG. 4 shows the flowchart of autonomous healthcare data exploration system **400**, which leverages multimode

health-trajectory predictors along with a consumer-centric database **401**. Autonomous healthcare data exploration system **400** includes the following components:

[0107] Inputs

[0108] Consumer-centric database (CCDB) **401** consisting of membership, benefit-plan history, consumer-touch-point history, claims, self-reported, consumer behavior marketing, provider, and evidence-based medicine data

[0109] Autonomous knowledge database, which is empty in the beginning, but will be populated with new and iteratively refined knowledge

[0110] Processing

[0111] Projection **403**: This step creates multiple projections of CCDB **401** over time so that one has a complete view of all that's happening to each consumer conditioned upon slowly-changing lifestyle, behavior, and psychographic parameters.

[0112] Overlapped frame feature extraction **405**: From each time frame of each projection space, one extracts an appropriate number of summarization and dynamic features so that we can track their trajectories over time.

[0113] Multimode health-trajectory predictors **407**: Predictors **407** predict future states of one's health around disease progression, engagement, and impact.

[0114] Past-future dynamic clustering **409**: Clustering is performed on the vector space spanned by the current set of features and predicted attributes. In one embodiment of such a system, the current set of features encompasses the parameterization of current disease conditions, utilization of medical resources, and lifestyle/health behavior. Predicted attributes may include disease progression, the level of impactability, and future cost. The key idea is to cluster consumers based on both where they are today and where they are likely to transition to in the future.

[0115] Anomalous cluster detection and merging **411**: Within each homogeneous cluster, one looks for outliers in joint and marginal spaces. Depending on the outlier-population size derived from each cluster, one merges outliers from multiple similar clusters to improve statistical power and significance.

[0116] Outcomes analysis **413**: For each outlier cluster, one looks for attributes with commonality and differences between outliers and normal cases. This search for common and uncommon attributes facilitates case-controlled outcomes analysis with drill-down along with the understanding of factors responsible for differences in outcomes.

[0117] Causal pathway analysis **415**: For each anomaly case identified, one uses a structural learning algorithm to induce a Bayesian network structure. Next, one ensures that causal parameters between control and test groups move in a logical way.

[0118] GUI for human confirmation **419**: Each discovered knowledge is presented to a human expert for final confirmation and inclusion into the autonomous knowledge discovery database **417**.

[0119] Outputs provided by autonomous healthcare data exploration system **400** include:

[0120] Extracted knowledge

Intelligent Health Benefit Design System

[0121] FIG. 5 shows an illustrative conceptual example of the optimal health benefit design in accordance with an embodiment of the invention. An intelligent benefit design system leverages ideas from consumer-understanding technologies, predictive modeling, impact analysis, and multi-objective optimization to design an individually tailored benefit product that balances the conflicting needs of moral hazard and social insurance by finding the acceptable ratio of profitability to subsidization for each product or plan configuration in a product bundle.

[0122] Element **515** in FIG. 5 shows a simplified two-dimensional efficient frontier in the two-dimensional space of premium and out-of-pocket (OOP) cost with an indifference curve. That is, higher premiums are generally associated with lower OOP costs and vice versa. An insurance company starts out with an initial set of product bundles **501**. If the company introduces a new product for which no prior enrollment data is available, then the product enrollment distribution is estimated using adaptive conjoint learning and prediction **503**. On the other hand, if product changes are evolutionary, then one can use prior enrollment data to develop and deploy predictive models to estimate the new product enrollment distribution given an initial set of product attributes **505**. As part of designing an adaptive conjoint analysis (ACA) questionnaire, one leverages consumer marketing database or demographic database from the U.S. Census Bureau so that the questionnaire can be tailored to each consumer **507**, **509**.

[0123] The fundamental idea is to iterate the process of adjusting product attributes, estimating product enrollment distributions, and calculating economic parameters (projected profit/loss as well as the level of subsidization inherent in a medical insurance product) of each product bundle so that we achieve an acceptable trade off between social insurance and moral hazard. That is, while the young and healthy are supposed to subsidize the cost of insurance for the old and sick, there needs to be an element of personal responsibility in benefit design so that people with poor health habits and beneficiary mentality do not abuse the entire healthcare system to the detriment of all **511**, **513**. In short, benefit design must deal effectively with risk factors that can be mitigated within socially acceptable means. The plot labeled **517** shows the relationship between individual prediction accuracy measured in R-sq or R² and group prediction accuracy measured in predictive ratio (PR) mean (μ) and standard deviation (σ). Individual predictive accuracy becomes less important as group size increases as in employer or group underwriting. However, in clinical settings and predicting benefit enrollment, where adverse selection can occur frequently, individual predictive accuracy is of paramount importance.

[0124] In healthcare, benefit design, according to prior art, is typically carried out by linking historical utilization and cost data to various benefit parameters, such as co-pay, deductible, co-insurance, maximum out-of-pocket, limits on Health Savings Account/Flexible Spending Account (HSA/FSA), etc. Then a loading factor (margin) is computed for each plan design, which sets the premium for the plan. Depending on the premium differential between plans, subsidization factors are calculated such that a plan attractive to predominantly the healthy (high-deductible plans) may sub-

sidize the cost of another plan that appeals primarily to the sick so that the concept of social insurance can be preserved in plan design.

[0125] An important consideration in benefit design is risk management. If benefit parameters are particularly attractive to a certain segment of population whose medical needs differ significantly from those of the general population, then such a plan has a high likelihood of attracting a biased population, which can lead to unexpected profit or loss depending on the direction of the bias. Unfortunately for health insurance companies, this phenomenon of biased population (called anti- or adverse selection) is not uncommon. The result is a cookie-cutter benefit design with a small number of selections so that the law of large numbers dominates the field.

[0126] More recently under the banner of consumer-directed health plan (CDHP), many payers started introducing high-deductible, low-premium plans. The theory of the case for CDHP is that high-deductible plans with some form of medical savings account will turn beneficiary-mentality patients into sophisticated healthcare consumers. Unlike other consumer industries, healthcare consumers may have hard time correlating actual high-quality care with a perceived one of at least based on RAND's quality metrics. Furthermore, the initial thrust of CDHP was to attract the cream-of-the-crop population from employers offering plans from multiple payers. That is, nimble new-to-the-market payers introduced CDHP products to employers desperate to cut soaring health benefit costs. The end result was that dinosaur payers were saddled with the undesirable segment of the population, hurting their bottom line.

[0127] Studies suggest that while the young and healthy are potential winners of CDHP, their opportunities for savings are limited because of restrictions in plan design, such as portability and investment. Results of post-CDHP health-resource utilizations and costs suggest mixed results with no clear trend. Perhaps mixed results are not surprising given the ambiguity of the theory of the case.

[0128] Perhaps the biggest shortcoming of the current health plan design is that few incorporate innovative design parameters, such as consumer-engagement strategies, incentives for lifestyle changes, and fun aspects in linking validated evidence-based-medicine guidelines, nutrition and exercise to health. Our design approach leverages the estimation of a consumer-preference function and projected utility functions derived from the impact analysis engine to move away from a cookie-cutter design and towards a tailored plan design that impacts health behavior change.

[0129] For new product launch **501**, one first proceeds with adaptive conjoint questionnaire (ACQ) **503** that is designed to minimize the number of questions leveraging predictive questionnaire construction. From ACQ responses, one can estimate a consumer preference function at an individual level. From a pool of initial product bundles with preset features, one estimates the overall enrollment distribution for a group (i.e., an employer). From the overall enrollment distribution and the outputs of multimode health-trajectory predictors, one computes profit/loss for each product and generate a three-dimensional picture of profit/loss and compressed two-dimensional objectives (i.e., minimize premium and out-of-pocket or OOP expense) as shown in relationships **515** and **517**. This picture will provide visual insights to facilitate the understanding of Pareto-efficient design parameters, which can lead to the reconfiguration of product features. This process of enrollment prediction and product reconfiguration is

iterative until the incremental change in product-feature reconfiguration is below an acceptable threshold.

[0130] After the product launch, one starts with a fresh data set, which represents the actual product selection behavior by consumers. Unlike in conjoint analysis, one does not have information on exactly which products consumers traded off before making product-selection decisions. One has the following information on consumers and their product-selection behavior:

1. Demographics and behavior marketing (x_{demo} , x_{cbm})
2. Prior product selection (x_{pps}), which doesn't exist for new consumers
3. Current product selection

[0131] The task at hand is to estimate a revised consumer preference function using real data. Let y and w denote the product-selection behavior and product features, respectively. Then, the estimation task is as follows:

$$\hat{y} = f(x_{demo}, x_{cbm}, x_{pps}, w, D(w, y)),$$

where $D(w, y)$ is a distance function between w and y , and $f(\bullet)$ can be estimated using parametric or nonparametric learning algorithms. Any differences between the conjoint and real-data models are stored in a database for continuous model adaptation and learning. More complex design with incentives requires utility functions associated with incentives from the impact-analysis engine. After estimating the consumer-preference function, there is a secondary step of identifying intervention opportunities given the characteristics of consumers choosing each product bundle. Based on utility functions and the outputs of the multimode health-trajectory predictors, the remaining task is to design an incentive program within each product bundle that will encourage high-risk members to participate in the program.

[0132] FIG. 6 shows an example of Markov modeling of assessing a target of opportunity in accordance with an embodiment of the invention. Markov model **600** shows a disease progression related to diabetes. Markov model **600** shows the probability of transitioning from one disease state to another disease state based on whether the consumer obtains a prescribed treatment. Additionally, disease states may depend on observed behavioral/lifestyle factors including the attributes of the consumer. Attributes may include the category of life style (e.g., "coach-potato") and level of education of the consumer. The type of treatment and the efficacy of the treatment may depend on the consumer's attributes.

[0133] With state **601**, a consumer, who is a "couch-potato," is determined to be a pre-diabetic. As determined by intervention opportunity finder **103** (as shown in FIG. 1), there is a probability p_{1a} **609a** of the consumer becoming a diabetic (state **603**) without any treatment and a probability p_1 **609b** if the consumer received a prescribed treatment (treatment_1). For example, EBM supercharger **300** may determine that the consumer can substantially reduce the probability of becoming a diabetic with a proper diet and exercise regime under the supervision of a dietician and/or exercise coach.

[0134] When the consumer becomes a diabetic, there is a probability of developing coronary artery disease (corresponding to state **605**). The corresponding treatment_2 (as determined by EBM supercharger) may be more radical than treatments. For example, treatment_1 may include one or more prescribed drugs that are typically more costly than

providing a dietician and/or exercise coach. (In general, as a disease progresses, the associated costs increase.) The probability of a diabetic developing coronary arterial disease without treatment is p_{2a} **611a** and p_2 **611b** with treatment.

[0135] In accordance with Markov model **600**, once a consumer has developed coronary arterial disease, the consumer may further develop renal failure and/or congestive heart failure (state **607**). The probability of developing renal failure/congestive failure is p_{3a} **113a** without treatment is and p_3 **611b** with treatment.

[0136] Markov model **600** may include states based on different attributes of a consumer. For example, state **615** is associated with the consumer having a physically active life style. Consequently, the transition probability of disease progression is typically smaller than a consumer having has a sedentary lifestyle (corresponding to state **601**, in which a consumer is classified as a “coach-potato”).

Exemplary Scenario

[0137] Sarah is a 45-year-old mother of two children, overweight, pre-diabetic, being treated for hypertension and hyperlipidemia. At work, she needs to enroll in a health benefit plan since her employer switched to a new payer, Global Health. In accordance with an embodiment of the invention, the following scenario that a consumer experiences.

[0138] Enrollment: Sarah is first given a combination of Predictive Health Risk Assessment (PHRA) interspersed with Adaptive Conjoint Analysis (ACA) questions. Even without single claims, PHRA calculates future health trajectories and guides Sarah through the benefit selection process based on an adaptive questionnaire tree designed to minimize the number of questions while maximizing predictive accuracy. She ends up selecting an HMO plan with various incentives for staying healthy. Impact analysis engine provided ROI's associated with incentives for consumers who fit Sarah's profile. She is given an instant analysis of her current health, likely health trajectories, and what she can do to prevent unpleasant outcomes. An interactive goal setting wraps up her first-day consumer experience with GH. Health trajectory predictors are based on PHRA/ACA questions, in which the optimal benefit design is part of resource allocation management (RAM). (With prior art, Sarah is typically given a list of traditional HMO, PPO, and Indemnity plans with a limited number of choices in deductibles, co-pays, and premium with health savings accounts.)

[0139] At-risk member identification: By virtue of PHRA, Sarah has already been identified as an at-risk member who can benefit from intervention. PHRA lists diabetes as a major risk factor given her current conditions, BMI, and lifestyle parameters inferred from external consumer behavior data obtained from Experian for a specific purpose of improving health guidance, not premium setting. Given her status, she gets a VAT call tailored to her situation, along with a two-page feedback/action plan letter based on her responses to the PHRA questionnaire all during the first week as part of a welcoming package. The Integrated Health Management Platform supports this function with health trajectory predictors, intervention opportunity finder, and RAM. (With prior art, since Sarah is a new member, GH must wait for claims data to accumulate before running

a predictive model that predicts 12-month future cost. Because of claims lag, the typical wait time is 6 months.)

[0140] Maintenance: Based on earlier communications, Sarah understands what to do. She takes PHRA frequently to report her progress and to see if her health scores are improving. Upon meeting her first goal of losing 10 lbs in 4 weeks and improving her health scores by 10%, GH sends her a USB pedometer. Now she uses it to keep track of her activity level daily, uploading to her personal Web portal at GH activity data, which provides additional data points to the IHM Platform in order to improve guidance for Sarah. Meanwhile the IHM Platform is exploring healthcare database autonomously, looking for patterns that precede low-to-high or high-to-low transitions so that it can update its knowledge database. Furthermore, it is constantly monitoring the relationship between intervention and outcomes to ensure that every member gets the best possible touch points to maximize population health using both high-tech and human interventions. The multimode health-trajectory predictors perform predictions both on a regular basis and asynchronously (event-driven). All IHM components work seamlessly to make this happen. (With prior art, not knowing the full extent of her risk factors, she may live her life as she normally does. One day, she feels chest pain and goes to ER. Upon examination, they find out that she needs heart bypass. Further blood test shows her blood glucose level at 175 mg/dl, which makes her a diabetic, further complicating her recovery. About 3 months after her bypass surgery, GH finally has her claims data in an electronic data warehouse. The indigenous PM now flags her as a high-risk member—a clear case of regression to the mean and fixing the door after a cow has already left. A nurse calls her to inquire if anything can be done to help her.)

Computer Implementation

[0141] FIG. 7 shows computer system **1** that supports an integrated health management platform (e.g., IHM platform **100** as shown in FIG. 1) in accordance with an embodiment of the invention. Elements of the present invention may be implemented with computer systems, such as the system **1**. Computer system **1** includes a central processor **10**, a system memory **12** and a system bus **14** that couples various system components including the system memory **12** to the central processor unit **10**. System bus **14** may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. The structure of system memory **12** is well known to those skilled in the art and may include a basic input/output system (BIOS) stored in a read only memory (ROM) and one or more program modules such as operating systems, application programs and program data stored in random access memory (RAM).

[0142] Computer **1** may also include a variety of interface units and drives for reading and writing data. In particular, computer **1** includes a hard disk interface **16** and a removable memory interface **20** respectively coupling a hard disk drive **18** and a removable memory drive **22** to system bus **14**. Examples of removable memory drives include magnetic disk drives and optical disk drives. The drives and their associated computer-readable media, such as a floppy disk **24** provide nonvolatile storage of computer readable instructions, data structures, program modules and other data for

computer 1. A single hard disk drive 18 and a single removable memory drive 22 are shown for illustration purposes only and with the understanding that computer 1 may include several of such drives. Furthermore, computer 1 may include drives for interfacing with other types of computer readable media.

[0143] A user can interact with computer 1 with a variety of input devices. FIG. 7 shows a serial port interface 26 coupling a keyboard 28 and a pointing device 30 to system bus 14. Pointing device 28 may be implemented with a mouse, track ball, pen device, or similar device. Of course one or more other input devices (not shown) such as a joystick, game pad, satellite dish, scanner, touch sensitive screen or the like may be connected to computer 1.

[0144] Computer 1 may include additional interfaces for connecting devices to system bus 14. FIG. 7 shows a universal serial bus (USB) interface 32 coupling a video or digital camera 34 to system bus 14. An IEEE 1394 interface 36 may be used to couple additional devices to computer 1. Furthermore, interface 36 may be configured to operate with particular manufacture interfaces such as FireWire developed by Apple Computer and i.Link developed by Sony. Input devices may also be coupled to system bus 14 through a parallel port, a game port, a PCI board or any other interface used to couple and input device to a computer.

[0145] Computer 1 also includes a video adapter 40 coupling a display device 42 to system bus 14. Display device 42 may include a cathode ray tube (CRT), liquid crystal display (LCD), field emission display (FED), plasma display or any other device that produces an image that is viewable by the user. Additional output devices, such as a printing device (not shown), may be connected to computer 1.

[0146] Sound can be recorded and reproduced with a microphone 44 and a speaker 66. A sound card 48 may be used to couple microphone 44 and speaker 46 to system bus 14. One skilled in the art will appreciate that the device connections shown in FIG. 7 are for illustration purposes only and that several of the peripheral devices could be coupled to system bus 14 via alternative interfaces. For example, video camera 34 could be connected to IEEE 1394 interface 36 and pointing device 30 could be connected to USB interface 32.

[0147] Computer 1 can operate in a networked environment using logical connections to one or more remote computers or other devices, such as a server, a router, a network personal computer, a peer device or other common network node, a wireless telephone or wireless personal digital assistant. Computer 1 includes a network interface 50 that couples system bus 14 to a local area network (LAN) 52. Networking environments are commonplace in offices, enterprise-wide computer networks and home computer systems.

[0148] A wide area network (WAN) 54, such as the Internet, can also be accessed by computer 1. FIG. 7 shows a modem unit 56 connected to serial port interface 26 and to WAN 54. Modem unit 56 may be located within or external to computer 1 and may be any type of conventional modem such as a cable modem or a satellite modem. LAN 52 may also be used to connect to WAN 54. FIG. 7 shows a router 58 that may connect LAN 52 to WAN 54 in a conventional manner.

[0149] It will be appreciated that the network connections shown are exemplary and other ways of establishing a communications link between the computers can be used. The existence of any of various well-known protocols, such as TCP/IP, Frame Relay, Ethernet, FTP, HTTP and the like, is presumed, and computer 1 can be operated in a client-server

configuration to permit a user to retrieve web pages from a web-based server. Furthermore, any of various conventional web browsers can be used to display and manipulate data on web pages.

[0150] The operation of computer 1 can be controlled by a variety of different program modules. Examples of program modules are routines, programs, objects, components, and data structures that perform particular tasks or implement particular abstract data types. The present invention may also be practiced with other computer system configurations, including hand-held devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, network PCS, minicomputers, mainframe computers, personal digital assistants and the like. Furthermore, the invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0151] In an embodiment of the invention, central processor unit 10 determines health trajectory predictors from HRA data 151, claims data 153, and CBM data 155 (as shown in FIG. 1), which are obtained through LAN 152 and WAN 154. Central processor unit 10 may also provide the functionalities of intervention opportunity finder 103, resource allocation manager 105, and impact analysis engine 107. Consequently, central processor unit 10 may provide a target of opportunity for a consumer from evidence-based medicine (EBM) guidelines or medical journals 351 (as shown in FIG. 351). EBM guidelines (corresponding to EBM database 317) and electronic health records (EHR) (corresponding to EHR database 203) may be retrieved from hard disk drive 18.

[0152] As can be appreciated by one skilled in the art, a computer system with an associated computer-readable medium containing instructions for controlling the computer system may be utilized to implement the exemplary embodiments that are disclosed herein. The computer system may include at least one computer such as a microprocessor, a cluster of microprocessors, a mainframe, and networked workstations.

[0153] While the invention has been described with respect to specific examples including presently preferred modes of carrying out the invention, those skilled in the art will appreciate that there are numerous variations and permutations of the above described systems and techniques that fall within the spirit and scope of the invention as set forth in the appended claims.

I claim:

1. An apparatus for processing medical literature comprising:
 - a text mining module detecting an outcome study for a medical treatment from a medical publication;
 - a rule induction module mapping the medical treatment to a diagnostic and procedural code;
 - a data analyzer configured to perform:
 - (a) accessing a database for health data using the diagnostic and procedural code, the health data spanning previous treatment for a population of consumers;
 - (b) associating an outcome metric for the medical treatment with a consumer group; and
 - (c) repeating (a) and (b) to determine another outcome metric for another consumer group; and

- a utility generator generating a utility function from the plurality of outcome metrics, the utility function gauging an efficacy of at least one intervention channel for a consumer.
2. The apparatus of claim 1, the text mining module detecting another outcome study for another medical treatment.
3. The apparatus of claim 1, the data analyzer determining the consumer group by clustering a subset of the population from at least one attribute of the population to form a cluster.
4. The apparatus of claim 3, the utility generator determining the utility function that relates a utility score for a selected intervention channel as applied to the cluster.
5. The apparatus of claim 4, the utility generator determining another utility score for another intervention channel as applied to another cluster.
6. The apparatus of claim 1, further comprising:
a confirmation interface presenting an abstract of the medical publication to a user and receiving, from the user, a notification that includes a confirmation of the abstract.
7. The apparatus of claim 6, the confirmation interface editing the abstract that is presented to the user.
8. The apparatus of claim 3, the data analyzer comparing a cluster dispersion measure of the cluster to a overall dispersion measure of the population and accepting the cluster from a ratio of the overall dispersion measure to the cluster dispersion measure.
9. The apparatus of claim 8, the data analyzer comparing the cluster dispersion measure and overall dispersion measure for each outcome distribution.
10. A method for processing medical literature comprising:
(a) detecting an outcome study for a medical treatment from a medical publication;
(b) mapping the medical treatment to a diagnostic and procedural code;
(c) accessing a database for health data using the diagnostic and procedural code, the health data spanning previous treatment for a population of consumers;
(d) associating an outcome metric for the medical treatment with a consumer group; and
(e) repeating (c) and (d) to determine another outcome metric for another consumer group; and
(f) generating a utility function from the plurality of outcome metrics, the utility function gauging an efficacy of at least one intervention channel for a consumer.
11. The method of claim 10, further comprising:
(g) detecting another outcome study for another medical treatment.
12. The method of claim 10, further comprising:
(g) determining the consumer group by clustering a subset of the population from at least one attribute of the population to form a cluster.

13. The method of claim 12, further comprising:
(h) determining the utility function that relates a utility score for a selected intervention channel as applied to the cluster.
14. The method of claim 13, further comprising:
(i) determining another utility score for another intervention channel as applied to another cluster.
15. The method of claim 10, further comprising:
(g) presenting an abstract of the medical publication to a user receiving, from the user, a notification that includes a confirmation of the abstract.
16. The method of claim 15, further comprising:
(h) editing the abstract that is presented to the user.
17. The method of claim 12, further comprising:
(h) comparing a cluster dispersion measure of the cluster to a overall dispersion measure of the population; and
(i) accepting the cluster from a ratio of the overall dispersion measure to the cluster dispersion measure.
18. The method of claim 17, further comprising:
(j) comparing the cluster dispersion measure and overall dispersion measure for each outcome distribution.
19. A computer-readable medium having computer-executable instructions to perform:
(a) detecting an outcome study for a medical treatment from a medical publication;
(b) mapping the medical treatment to a diagnostic and procedural code;
(c) accessing a database for health data using the diagnostic and procedural code, the health data spanning previous treatment for a population of consumers;
(d) associating an outcome metric for the medical treatment with a consumer group; and
(e) repeating (c) and (d) to determine another outcome metric for another consumer group; and
(f) generating a utility function from the plurality of outcome metrics, the utility function gauging an efficacy of at least one intervention channel for a consumer.
20. The computer-readable medium of claim 19, further configured to perform:
(g) determining the consumer group by clustering a subset of the population from at least one attribute of the population to form a cluster.
21. The computer-readable medium of claim 20, further configured to perform:
(h) determining the utility function that relates a utility score for a selected intervention channel as applied to the cluster.
22. The computer-readable medium of claim 20, further configured to perform:
(h) comparing a cluster dispersion measure of the cluster to a overall dispersion measure of the population; and
(i) accepting the cluster from a ratio of the overall dispersion measure to the cluster dispersion measure.

* * * * *