(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
19 December 2019 (19.12.2019)

WIPO | PCT

(10) International Publication Number
WO 2019/237193 A1

(51) International Patent Classification:
G06F 16/906 (2019.01) *G16H 50/70* (2018.01)
G06F 16/80 (2019.01) *G06N 3/02* (2006.01)
G06F 16/904 (2006.01) *G06N 3/12* (2006.01)
G06N 20/00 (2019.01)

(21) International Application Number:
PCT/CA2019/050823

(22) International Filing Date:
11 June 2019 (11.06.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
3007786 11 June 2018 (11.06.2018) CA
16/005,069 11 June 2018 (11.06.2018) US

(71) Applicant: **ODAIA INTELLIGENCE INC.** [CA/CA]; 1 University Avenue, 3rd Floor, Toronto, Ontario M5J 2P1 (CA).

(72) Inventor: **ANDRITSOS, Periklis**; 439 Soudan Avenue, Toronto, Ontario M4S 1X1 (CA).

(74) Agent: **NORTON ROSE FULBRIGHT CANADA LLP**; 222 Bay Street, Suite 3000, P.O. Box 53, Toronto, Ontario M5K 1E7 (CA).

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

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,

(54) **Title:** DATA VISUALIZATION PLATFORM FOR EVENT-BASED BEHAVIOR CLUSTERING

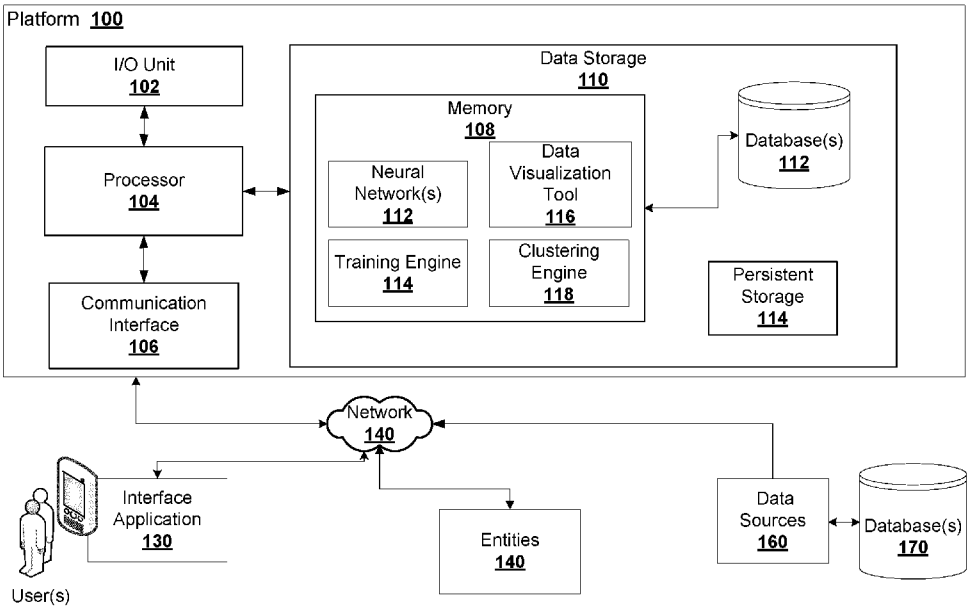

FIG. 1

(57) **Abstract:** A platform for processing event traces to generate clusters of journey maps for an interactive user display. The clustering can be implemented using a genetic process. The clustering can be implemented using a distance measures. Other clustering techniques can be used. The platform can use the clustering results to generate interactive dynamic visualizations with interactive selectable portions. The platform can integrate processing mining and journey maps to generate interactive dynamic visualizations.

WO 2019/237193 A1

[Continued on next page]

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

Published:

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

— *in black and white; the international application as filed contained color or greyscale and is available for download from PATENTSCOPE*

TITLE: DATA VISUALIZATION PLATFORM FOR EVENT-BASED BEHAVIOR CLUSTERING

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to Canadian Patent Application No. 3,007,786 filed June 11, 2018 and U.S. Patent Application No. 16/005,069 filed June 11, 2018.

FIELD

[0002] The present disclosure generally relates to the field of data visualization, machine learning, cluster analysis, and computer interfaces.

INTRODUCTION

[0003] Embodiments described herein relate to a data visualization for event-based behavior analysis or trace analysis. Machine learning is a field of computer science that involves programming code configured to learn, adapt or improve performance on a specific task with data, without being explicitly programmed. Cluster analysis can involve code that processes input data to group a set of objects in clusters based on detected patterns or similarities. That is, objects grouped in a cluster are more similar (in some sense) to each other than to those in other clusters. Machine learning can involve cluster analysis to learn patterns and similarities when processing input data. Machine learning and cluster analysis can generate output data that can be used to create dynamic, interactive visual elements for computer interfaces. The computer interfaces can receive control commands that can be received as input data and trigger generation of updated visual elements.

SUMMARY

[0004] In accordance with an aspect, there is provided a platform for processing event traces to generate clusters of the event traces for an interactive user display.

[0005] In various further aspects, the disclosure provides corresponding systems and devices, and logic structures such as machine-executable coded instruction sets for implementing such systems, devices, and methods.

[0006] In accordance with an aspect, there is provided a platform with data storage device having a data warehouse model for storing event traces, each event trace having attributes that

indicate activities over time. The platform has a processor configured to process machine executable instructions to generate visual elements for an interactive interface application by: generating hierarchical cluster for the event traces by grouping event traces having similar attributes; generating a set of representative attributes for a set of event traces of the hierarchical cluster, the set of representative attributes computed based on the similar attributes; generating the visual elements for multiple views for the interface application, the multiple views having a first view indicating a pattern of activities for the set of representative attributes; a second view indicating the hierarchical cluster and the pattern of activities within the hierarchical cluster; and a third view indicating descriptors for the hierarchical cluster, the multiple views having a plurality of selectable portions; controlling rendering of the interface application at a device to display the multiple views synchronously; and responsive to a selection of a selectable portion of the plurality of selectable portions, controlling rendering of the interface application at the device to update the multiple views synchronously based on the selected portion to navigate the hierarchical cluster.

[0007] In some embodiments, the processor generates the hierarchical cluster by computing a distance measure for the attributes of the event traces to compute the similar attributes.

[0008] In some embodiments, the processor generates the hierarchical cluster by computing a distance measure based on a Levenshtein distance.

[0009] In some embodiments, the processor generates the hierarchical cluster by computing a distance between clusters of the hierarchical cluster.

[0010] In some embodiments, the pattern of activities indicates the set of representative attributes for a segment of the hierarchical cluster that corresponds to the set of event traces.

[0011] In some embodiments, the second view indicates the segment of the hierarchical cluster for the pattern of activities indicated in the first view.

[0012] In some embodiments, the selection portion is within a view of the multiple views to trigger an update to the other views of the multiple views.

[0013] In some embodiments, the processor updates the multiple views synchronously based on the selected portion to navigate the hierarchical cluster at different levels of abstraction or granularity.

[0014] In some embodiments, the processor generates the hierarchical cluster based on a proximity measure using an order of activities for the event traces.

[0015] In some embodiments, the processor computes salient characteristics for the descriptors to indicate indexes based on a chi-square test.

[0016] In some embodiments, the processor receives a navigation goal at the interface application, the goal indicating an attribute and responsive to the navigation goal, controlling rendering of the interface application at the device to update the multiple views synchronously to indicate data based on the navigation goal.

[0017] In some embodiments, the processor generates the hierarchical cluster based on layers, a layer corresponding to a number of event traces that will be grouped based on the similar attributes, the layers corresponding to a height of the hierarchical cluster.

[0018] In some embodiments, the processor generates the set of representative attributes based on the similar attributes and a frequent sequences mining process.

[0019] In accordance with an aspect, there is provided a non-transitory computer readable medium storing machine executable instructions to configure a processor to: generate hierarchical cluster for the event traces by grouping event traces having similar attributes, each event trace having attributes that indicate activities over time; generate a set of representative attributes for a set of event traces of the hierarchical cluster, the set of representative attributes computed based on the similar attributes; generate the visual elements for multiple views for the interface application, the multiple views having a first view indicating a pattern of activities for the set of representative attributes; a second view indicating the hierarchical cluster and the pattern of activities within the hierarchical cluster; and a third view indicating descriptors for the hierarchical cluster, the multiple views having a plurality of selectable portions; control rendering of the interface application at a device to display the multiple views synchronously; and responsive to a selection of a selectable portion of the plurality of selectable portions, control rendering of the interface application at the device to update the multiple views synchronously based on the selected portion to navigate the hierarchical cluster.

[0020] In some embodiments, the machine executable instructions configure the processor to generate the hierarchical cluster by computing a distance measure for the attributes of the event traces to compute the similar attributes.

[0021] In some embodiments, the machine executable instructions configure the processor to generate the hierarchical cluster by computing a distance between clusters of the hierarchical cluster.

[0022] In some embodiments, the machine executable instructions configure the processor to update the multiple views synchronously based on the selected portion to navigate the hierarchical cluster at different levels of abstraction or granularity.

[0023] In some embodiments, the machine executable instructions configure the processor to receive a navigation goal at the interface application, the goal indicating an attribute and responsive to the navigation goal, controlling rendering of the interface application at the device to update the multiple views synchronously to indicate data based on the navigation goal.

[0024] In some embodiments, the machine executable instructions configure the processor to generate the hierarchical cluster based on layers, a layer corresponding to a number of event traces that will be grouped based on the similar attributes, the layers corresponding to a height of the hierarchical cluster.

[0025] In accordance with an aspect, there is provided a computer process to generate visual elements for an interactive interface application. The process involves, at a processor, generating hierarchical cluster for the event traces by grouping event traces based on distance measures, each event trace having attributes that indicate activities over time, generating a set of representative attributes for a set of event traces of the hierarchical cluster, the set of representative attributes computed based on the grouped event traces; generating the visual elements for multiple views for the interface application, the multiple views having a first view indicating a pattern of activities for the set of representative attributes; a second view indicating the hierarchical cluster and the pattern of activities within the hierarchical cluster ; and a third view indicating descriptors for the hierarchical cluster, the multiple views having a plurality of selectable portions; controlling rendering of the interface application at a device to display the multiple views synchronously; and responsive to a selection of a selectable portion of the plurality of selectable portions, controlling rendering of the interface application at the device to update the multiple views synchronously based on the selected portion, wherein the selection portion is within a view of the multiple views to trigger an update to the other views of the multiple views.

[0026] In accordance with an aspect, there is provided a platform with a data storage device for storing event traces, the event traces having attributes that indicate activities over time. The

platform has a processor configured to process machine executable instructions to generate visual elements for an interactive interface application by: computing clusters of the event traces based on a genetic process function, each cluster corresponding to a set of event traces and a representative event trace based on representative attributes of the set of event traces, the genetic process function mapping each event trace to a cluster; generating the visual elements for an interface application, the visual elements indicating the clusters and, for each cluster, the representative event trace based on the representative attributes, the representative event trace summarizing the set of event traces corresponding to the cluster; controlling rendering of the interface application at a device to display the visual elements and a plurality of selectable portions; and responsive to a selection of a selectable portion of the plurality of selectable portions, controlling rendering of the interface application at the device to navigate the clusters.

[0027] In some embodiments, the processor is configured to implement the genetic process function by: evaluating initial set representative event traces to generate elite set representative event traces, generating additional initial representative event traces using a transformation process, evaluating the additional initial representative event traces to generate additional elite set representative event traces, continuing the generating and the evaluating until stopping criterion is met, using resulting elite representative event traces as the representative event traces for the clusters.

[0028] In some embodiments, for each event trace, the attributes are represented as a two-dimensional tuple of activities over time and contextual data and the genetic process function can evaluate the two-dimensional tuples of the event traces.

[0029] In some embodiments, the processor computes, for each cluster, the representative event trace as a pattern of activities to summarize the patterns contained within the set of event traces of the cluster.

[0030] In some embodiments, the processor implements the genetic process function by pre-processing the event traces, determining an initial population that corresponds to a number of the clusters, evaluating initial representative event traces, evaluating stop criteria, implementing a genetic operation to generate additional initial representative event traces, outputting a map abstraction for use in generating additional visual elements.

[0031] In some embodiments, the processor is configured to evaluate the initial representative event traces based on quality criteria for fitness, a number of representatives, a contextual distance and an average quality.

[0032] In some embodiments, the processor is configured to weigh the quality criteria to generate an overall quality and average the overall quality.

[0033] In some embodiments, the processor is configured to evaluate the initial representative event traces based on internal and external evaluation metrics.

[0034] In accordance with an aspect, there is provided a platform having a data storage device for storing event traces, the event traces having attributes that indicate activities over time, the event traces relating to current data and historic data. The platform has a processor configured to process machine executable instructions to generate visual elements for an interactive interface application by: generating a mapping of process mining elements to the event traces, the mapping having an XML-based set of concepts, the mapping linking process de facto models and de jure models to actual and expected customer experiences; generating a data structure for hierarchical components of a journey map, each component having an element and an attribute, the journey map having at least one element being a root node of the hierarchical components, the actual and expected customer experiences corresponding to the journey map; generating a visual elements for the components of the journey map; and transmitting the visual elements to an interface application for display on a device;

[0035] In accordance with an aspect, there is provided a platform having a data storage device for storing event traces, the event traces having attributes that indicate activities over time. The platform has a processor configured to process machine executable instructions to generate visual elements for a map abstractor of an interactive interface application by: generating a process tree for the event traces by grouping the event traces based on similar attributes, the process tree having nodes corresponding to the event traces, the process tree having leaf nodes; parsing the process tree; starting from leaf nodes of the parsed process tree, iteratively generating a prompt at an interface application to merge a set of event traces that belong to a subset of the parsed process tree; upon receiving confirmation to merge the set of event traces, generating a name for the set of event traces and replacing the nodes corresponding to the set of event traces with a new node indicating the name and a set of representative attributes for the set of event traces; generating the visual elements for an interface application, the visual elements indicating the process tree and the new nodes, the visual elements representing an abstracted

process tree; controlling rendering of the interface application at a device to display the visual elements.

[0036] In some embodiments, a process tree is an abstract hierarchical representation of a process model, where the leaves are annotated with activities and all the other nodes are annotated with operators.

[0037] In this respect, before explaining at least one embodiment in detail, it is to be understood that the embodiments are not limited in application to the details of construction and to the arrangements of the components set forth in the following description or illustrated in the drawings. Also, it is to be understood that the phraseology and terminology employed herein are for the purpose of description and should not be regarded as limiting.

[0038] Many further features and combinations thereof concerning embodiments described herein will appear to those skilled in the art following a reading of the instant disclosure.

DESCRIPTION OF THE FIGURES

[0039] Figure 1 is a schematic diagram of platform for clustering and data visualizations according to some embodiments.

[0040] Figure 2 shows a trace depicting factors influencing vaccination decisions.

[0041] Figure 3 shows an example trace model.

[0042] Figure 4 shows the mapping of the XES concepts to a model generated by platform.

[0043] Figure 5 shows the process mining framework extended for trace models.

[0044] Figure 6 shows the application of process mining activities on trace data and results.

[0045] Figure 7 shows an example hierarchical clustering that can be generated by the platform.

[0046] Figure 8 is an example screenshot visible at interface application and points out three views that are available to navigate the clusters.

[0047] Figure 9 shows a screenshot of interface application for goal setting impact on a tree.

[0048] Figure 10 illustrates a cluster implementation containing example steps.

[0049] Figure 11 illustrates a preview of the parameters when uploading a dataset.

[0050] Figure 12 is an example screenshot of interface application.

[0051] Figure 13 is a diagram that shows that the process mining framework as well as the input data are relevant in a customer journey analytics context.

[0052] Figure 14 shows that these business process model (BPM) and journey maps convey different information.

[0053] Figure 15 depicts phases of determining a representative journey.

[0054] Figure 16 provides an overview of an example genetic process.

[0055] Figure 17 depicts *Top* from the example given in Figure 16.

[0056] Figure 18 illustrates the evolution of the evaluation criteria.

[0057] Figure 19 illustrates a graph of the distribution.

[0058] Figure 20 illustrates a value frequency table.

[0059] Figure 21 illustrates how six journeys are generated from two generative journeys.

[0060] Figure 22 depicts a clustering process.

[0061] Figure 23 shows a solution that reduces the Jaccard distance.

[0062] Figure 24 shows a drop in the average absolute distance in the number of journeys.

[0063] Figure 25 shows results of the genetic algorithm.

[0064] Figure 26 shows results of the genetic algorithm.

[0065] Figure 27 shows results using different setting configurations.

[0066] Figure 28 depicts an example for the handling of reviews for a journal.

[0067] Figure 29 displays relations using a process tree.

[0068] Figure 30 depicts map abstractor steps.

[0069] Figure 31 depicts results of the process tree.

[0070] Figure 32 shows a screen capture of the application when merging the activities.

[0071] Figure 33 shows journey maps at three levels of abstraction.

DETAILED DESCRIPTION

[0072] Embodiments of methods, systems, and apparatus are described through reference to the drawings.

[0073] Figure 1 is a schematic diagram of platform 100 for clustering and data visualizations according to some embodiments. The platform 100 can implement aspects of the processes described herein for clustering data using clustering engine 118 and generating data visualizations using data visualization tool 116.

[0074] The platform 100 generates data visualizations for interface application 130 and implements event-based behavior analysis or trace analysis using clustering engine 118. The clustering engine 118 implements cluster analysis using code that processes input data to group a set of objects in clusters based on detected patterns or similarities. Machine learning involves programming code stored in memory 108 that is configured to learn, adapt or improve performance on a specific task with data, without being explicitly programmed using training engine 114 and neural networks 112. Clustering engine 118 can implement machine learning with cluster analysis to learn patterns and similarities when processing input data. Clustering engine 118 uses machine learning and cluster analysis to generate output data, and integrates with data visualization tool 116 to generate dynamic, interactive visual elements based on the output data for display at interface application 130. The interface application 130 can receive control commands that can be received as input data and trigger generation of updated visual elements by clustering engine 118 and data visualization tool 116.

[0075] The platform 100 has a processor 104 and a memory 108 storing machine executable instructions to configure the processor 104 to receive input data (from e.g. entities 140, data sources 160). The platform 100 connects to interface application 130, entities 140, network endpoints 150, and data sources 160 (with databases 170) using network 140. Entities 140, interface application 130, and data sources 160 (with databases 170) can interact with the platform 100 to provide input data and receive output data. For example, the platform 100 can Network 140 (or multiple networks) is capable of carrying data and can involve wired connections,

wireless connections, or a combination thereof. Network 140 may involve different network communication technologies, standards and protocols, for example. The interface application 130 can be installed on user device to display an interface of visual elements, for example.

[0076] The platform 100 can include an I/O Unit 102, a processor 104, communication interface 106, and data storage 110. The processor 104 can execute instructions in memory 108 to implement aspects of processes described herein. The processor 104 can execute instructions in memory 108 to configure neural networks 112, training engine 114, data visualization tool 116, clustering engine 118, and other functions described herein. For example, training engine 114 can train neural networks 112 using training data for cluster analysis. The clustering engine 118 can use the trained neural networks 112 to generate output data for data visualization tool 116. The neural networks 112 can generate predictions for event traces or journeys, for example. The clustering engine 118 may be software (e.g., code segments compiled into machine code), hardware, embedded firmware, or a combination of software and hardware, according to various embodiments.

[0077] The data storage device 110 has a data warehouse model for storing event traces. An event trace has attributes that indicate activities over time, for example. An event trace can be a set of touchpoints or a sequence of activities. The data storage device 110 can store event logs that represent a collection of event traces. The event traces can be used by platform 100 to generate journeys that can be actual journeys or representative journeys. The event traces can be based on real-time (e.g. current) and/or historical data. The data visualization tool 116 can generate visual elements for an interactive interface application 130 by generating clusters for the event traces using the clustering engine 118 to group event traces having similar attributes. The clustering engine 118 can cluster similar traces that are close together using distance measures. The clustering engine 118 can generate the hierarchical cluster by computing a distance measure for the attributes of the event traces to compute the similar attributes. The clustering engine 118 can generate the hierarchical cluster by computing a distance measure based on a Levenshtein distance. The clustering engine 118 can generate the hierarchical cluster by computing a distance between clusters of the hierarchical cluster. The clustering engine 118 can generate a set of representative attributes for a set of event traces of the hierarchical cluster. The set of representative attributes computed based on the similar attributes. Neural networks 112 can be used to represent the event traces and the attributes or to determine representative attributes based on patterns. Training engine 114 can train the neural networks 112 using event traces to identify patterns of attributes or activities. In some embodiments, the pattern of activities indicates

the set of representative attributes for a segment of the hierarchical cluster that corresponds to the set of event traces. The clustering engine 118 can generate the hierarchical cluster based on a proximity measure using an order of activities for the event traces. The clustering engine 118 can generate the hierarchical cluster based on layers. A layer can correspond to a number of event traces that will be grouped based on the similar attributes, the layers corresponding to a height of the hierarchical cluster. In some embodiments, the clustering engine 118 generates the set of representative attributes based on the similar attributes and a frequent sequences mining process.

[0078] The data visualization tool 116 can generate the visual elements for multiple views for the interface application. The multiple views have a first view indicating a pattern of activities for the set of representative attributes. A pattern entails many event traces and indicates representative attributes for the activities or activity it represents. The clustering engine 118 can use the neural network 112 to detect the patterns, for example. A second view can indicate the hierarchical cluster and the pattern of activities within the hierarchical cluster. The second view can highlight the location of the pattern of activities in the visual representation of the hierarchical cluster. In some embodiments, the second view indicates the segment of the hierarchical cluster for the pattern of activities indicated in the first view. A third view indicates descriptors for the hierarchical cluster. The third view corresponds to the level of granularity of the first and second view and the descriptors indicate why the subset of the cluster is interesting, for example. The descriptors can indicate distinct traces, total traces and salient characteristics. In some embodiments, the processor computes salient characteristics for the descriptors to indicate indexes based on a chi-square test.

[0079] The multiple views have selectable portions that trigger the data visualization tool 116 to generate updates for the visual elements. The platform 100 controls rendering of the interface application 130 at a device to display the multiple views synchronously. Responsive to a selection of a selectable portion of the plurality of selectable portions, the platform 100 controls rendering of the interface application 130 at the device to update the multiple views synchronously based on the selected portion to navigate the hierarchical cluster. For example, the selectable portions can select attributes (e.g. activities) to filter the views (e.g. event trace contains selected activity X). In some embodiments, the selection portion is within a view of the multiple views to trigger an update to the other views of the multiple views. In some embodiments, the data visualization tool 116 updates the multiple views synchronously based on the selected portion to navigate the hierarchical cluster at different levels of abstraction or granularity.

[0080] In some embodiments, the platform 100 receives a navigation goal at the interface application 130. The goal indicates an attribute, for example. Responsive to the navigation goal, the platform 100 controls rendering of the interface application 130 at the device to update the multiple views synchronously to indicate data based on the navigation goal.

[0081] In some embodiments, the interface application 130 can show a view indicating actual event traces and a view showing representative event traces. The representative event traces being less than the actual event traces to provide an abstracted representation of the event traces. It may be difficult to see trends and insights from actual events traces if there is too much raw data shown at the interface application 130.

[0082] In some embodiments, clustering engine 118 can compute clusters of the event traces based on a genetic process function. In some embodiments, for each event trace, the attributes are represented as a two-dimensional tuple of activities over time and contextual data and the genetic process function can evaluate the two-dimensional tuples of the event traces.

[0083] Each cluster can correspond to a set of event traces and a representative event trace based on representative attributes of the set of event traces. The genetic process function maps each event trace to a cluster. The data visualization tool 116 can generate the visual elements for the interface application 130 and the visual elements can indicate the clusters and, for each cluster, the representative event trace based on the representative attributes. The representative event trace summarizes the set of event traces corresponding to the cluster. The data visualization tool 116 can control rendering of the interface application 130 at the device to display the visual elements and a plurality of selectable portions. Responsive to a selection of a selectable portion of the plurality of selectable portions, data visualization tool 116 can control rendering of the interface application 130 to navigate the clusters. The clustering engine 118 can determine segments for the event traces and then evaluate them to determine the clusters, for example. The clustering engine 118 can group the event traces into X groups and then identify a representative event trace for each group. The representative event trace does not have to be an actual event trace or representative of actual raw data but can summarize the group or cluster of event traces it represents. In some embodiments, the platform 100 can use neural networks 112 for natural language processing of text data that can form part of event traces to determine a sentiment score or sentiment indicator. This sentiment score or sentiment indicator can be an attribute for the event trace, for example. The genetic approach goes through different

combinations of clusters and evaluates the clusters to determine the best segments of event traces.

[0084] In some embodiments, clustering engine 118 can implement the genetic process function by: evaluating initial set representative event traces to generate elite set representative event traces, generating additional initial representative event traces using a transformation process, evaluating the additional initial representative event traces to generate additional elite set representative event traces, continuing the generating and the evaluating until stopping criterion is met, using resulting elite representative event traces as the representative event traces for the clusters.

[0085] In some embodiments, clustering engine 118 computes, for each cluster, the representative event trace as a pattern of activities to summarize the patterns contained within the set of event traces of the cluster. In some embodiments, clustering engine 118 implements the genetic process function by pre-processing the event traces, determining an initial population that corresponds to a number of the clusters, evaluating initial representative event traces, evaluating stop criteria, implementing a genetic operation to generate additional initial representative event traces, outputting a map abstraction for use in generating additional visual elements. In some embodiments, clustering engine 118 is configured to evaluate the initial representative event traces based on quality criteria for fitness, a number of representatives, a contextual distance and an average quality. In some embodiments, the clustering engine 118 is configured to weigh the quality criteria to generate an overall quality and average the overall quality. In some embodiments, the clustering engine 118 is configured to evaluate the initial representative event traces based on internal and external evaluation metrics.

[0086] In some embodiments, the data visualization tool 116 is configured to generate visual elements for an interactive interface application by: generating a mapping of process mining elements to the event traces, the mapping having an XML-based set of concepts. The mapping links process de facto models and de jure models to actual and expected customer experiences; generating a data structure for hierarchical components of a journey map, each component has an element and an attribute. The mapping can be stored in data storage 110. The journey map has at least one element being a root node of the hierarchical components. The actual and expected customer experiences correspond to the journey map. The data visualization tool 116 is configured to generate visual elements for the components of the journey map and transmit the visual elements to an interface application for display on a device.

[0087] In some embodiments, the data visualization tool 116 is configured to generate visual elements for a map abstractor of an interactive interface application by: generating a process tree for the event traces by grouping the event traces based on similar attributes, the process tree having nodes corresponding to the event traces, the process tree having leaf nodes. The data visualization tool 116 can parse the process tree, and starting from leaf nodes of the parsed process tree, iteratively generates a prompt at the interface application 130 to merge a set of event traces that belong to a subset of the parsed process tree. Upon receiving confirmation to merge the set of event traces, data visualization tool 116 can generate a name for the set of event traces and replace the nodes corresponding to the set of event traces with a new node indicating the name and a set of representative attributes for the set of event traces. This updated representation can be stored at data storage 110. The data visualization tool 116 can generate the visual elements for the interface application 130 that indicate the process tree and the new nodes for representing an abstracted process tree and controls rendering of the interface application 130 to display the visual elements. In some embodiments, a process tree is an abstract hierarchical representation of a process model, where the leaves are annotated with activities and all the other nodes are annotated with operators. In some embodiments, a merge of activities for a first data set can enable training engine 114 and neural network 112 to learn a similar model with the same activities that can be applied the same merge or abstraction rules.

[0088] The I/O unit 102 can enable the platform 100 to interconnect with one or more input devices, such as a keyboard, mouse, camera, touch screen and a microphone, and/or with one or more output devices such as a display screen and a speaker.

[0089] The processor 104 can be, for example, any type of general-purpose microprocessor or microcontroller, a digital signal processing (DSP) processor, an integrated circuit, a field programmable gate array (FPGA), a reconfigurable processor, or any combination thereof.

[0090] Memory 108 may include a suitable combination of any type of computer memory that is located either internally or externally such as, for example, random-access memory (RAM), read-only memory (ROM), compact disc read-only memory (CDROM), electro-optical memory, magneto-optical memory, erasable programmable read-only memory (EPROM), and electrically-erasable programmable read-only memory (EEPROM), Ferroelectric RAM (FRAM) or the like. Data storage devices 110 can include memory 108, databases 112 (e.g. graph database), and persistent storage 114.

[0091] The communication interface 106 can enable the platform 100 to communicate with other components, to exchange data with other components, to access and connect to network resources, to serve applications, and perform other computing applications by connecting to a network (or multiple networks) capable of carrying data including the Internet, Ethernet, plain old telephone service (POTS) line, public switch telephone network (PSTN), integrated services digital network (ISDN), digital subscriber line (DSL), coaxial cable, fiber optics, satellite, mobile, wireless (e.g. Wi-Fi, WiMAX), SS7 signaling network, fixed line, local area network, wide area network, and others, including any combination of these.

[0092] The platform 100 can be operable to register and authenticate users (using a login, unique identifier, and password for example) prior to providing access to applications, a local network, network resources, other networks and network security devices. The platform 100 can connect to different machines or entities 140.

[0093] The data storage 110 may be configured to store information associated with or created by the platform 100. Storage 110 and/or persistent storage 114 may be provided using various types of storage technologies, such as solid state drives, hard disk drives, flash memory, and may be stored in various formats, such as relational databases, non-relational databases, flat files, spreadsheets, extended markup files, and so on.

[0094] The platform 100 can implement a processing mining based model for trace analysis or event-based behavior analysis. The platform 100 can generate an appropriate mapping for process mining in order to help trace user data and customer experience. The platform 100 can use XML-based set of concepts to implement the above mapping, for example. The platform 100 can codify a parallel that exists between different types of business process models (e.g. “de facto” and “the jure”) and the ways that exist to analyze a customer experience (e.g. actual and expected).

[0095] Customer journey maps (journey maps) are data structures that model data relating to customers’ behavior. Platform 100 interacts with interface application 130 to display visual representations of journey maps generated by data visualization tool 116 to assist in decision making. The platform 100 can use a journey map model with process mining, a data analytics technique to assess the impact of the journey’s duration on the customer experience.

[0096] Customer journey map is a technique that enables professionals to better understand customers’ experiences when they interact with the steps involved in a service. These

interactions, called touchpoints, are increasing and can be used by customers in erratic ways. Hence, being able to use journey maps is becoming increasingly important for companies. journey maps can be applied to sales, healthcare, and library science, for example.

[0097] Figure 2 shows a journey map 200 depicting factors influencing vaccination decisions. The diagram indicates elements that can be found on journey maps: touchpoints (e.g., consultation experience), stages (e.g., pre-vaccination), and emotions (i.e., factors influencing the decision).

[0098] Two example journey maps include anticipated journey (e.g. expected journey) and actual journey, which aims to describe how the journey was "really" experienced by customers. There can be an interplay between expected and actual journeys. For instance, traces of customer journeys available in platform 100 and/or data sources 160 can be used to build a journey map from data or facts. Then, this journey map can be compared with an expected journey map generated by a prediction model.

[0099] Process mining relates to the use of process models and events logs to deliver fact-based insights. The platform 100 can use process mining for journey maps. Process mining works with event logs, a sequential format ideal for representing journey maps. Working with expected and actual models is at the core of a process mining framework. An example relates to process mining in the context of sales and how a structured approach can facilitate a sales method. Embodiments can apply process to journey map data.

[00100] Embodiments can define components of journey maps, develop a model for journey maps, use process mining to analyze journey maps, and additional analytics-oriented usage of journey maps.

[00101] There are various components of a journey map such as customer, journey, mapping, goal, touchpoint, timeline, channel, stage, experience, lens, and multimedia, for example.

[00102] Customer. A customer is the stakeholder experiencing a service. This can include people such as patients, students, or software users. A customer can be defined by sociodemographic information. When a customer is mentioned as a fictional character, the term "persona" is sometimes used.

[00103] Journey. A journey map contains at least one journey, which is a path through a graph of nodes that relates to a path followed by a customer. An example journey map can relate to

internal stakeholders to describe what an ideal journey would look like, which identifies opportunities for novel services or is employed as a diagnostic tool. This can be referred to as an expected journey. In contrast, the actual journey showcases how a journey is experienced by the customer, finds existing customers' problems or needs or pictures the consumption of services by customers.

[00104] Mapping. Mapping is a process consisting of tracking and describing customers' responses and experiences when using a service. Ultimately, these elements are reported on a map.

[00105] Goal. A customer journey should be mapped with a goal in mind, which is also referred to as scenario, prompts, story, or main intention. It triggers interactions with users, and streamlines the thought process for users.

[00106] Touchpoint. A touchpoint is an interaction between customers and companies' products or services such as "searching for a product", or "finding seats". The arrangement of touchpoints can be cyclic: a customer can iterate a few times over the same touchpoints. Moreover, the arrangement can be non-linear: (1) most of the time, the customer will not go through all the existing touchpoints; (2) the customer might miss a planned touchpoint; and (3) the customer can unexpectedly quit the journey.

[00107] Timeline. The timeline describes the duration of the journey from the first until the last touchpoint. Data can also not have a timestamp. A number attached to an event (i.e., touchpoint) can depict the sequence within the timeline.

[00108] Channel. The channel is the method chosen by the customer to interact with the touchpoint such as a "reference desk" or "social media".

[00109] Stage. A stage encompasses several touchpoints. Some authors used the splits: before, during, and after the experience, but employing domain-related steps is also possible. For instance, the stage refers to the waterfall model (i.e., software development). Some journey maps might not use stages.

[00110] Experience. The experience encompasses customers' feedback and emotions. There can be elements to express the experience. One element can be the emotion. Using only one continuum of emotions–such as unhappy to happy–may fail to depict a customer's experience. Thus, describing the emotion can require some flexibility or a range of values. An element can be

a scale that measures how positive or negative the experience was for the customer. Another element may be customers' quotes to represent what customers have been through.

[00111] Lens. Some components of journey maps are domain-specific. For instance, there can be an appended layer to indicate the weather because it impacts customer satisfaction when using the service. The layer can refer to the term lens to reflect that multiple views are possible on the same map. Suggestions and opportunities are some other examples of lenses superposed on top of touchpoints. They can promote reflection and analysis of what happened during the journey. The lens can trigger different visual representations on interface application 130 and data visualization tool 116 can generate different lens for a journey map involving different visual elements.

[00112] Multimedia. The usage of multimedia as part of a visual representation on interface application 130 can make a journey map engaging. For instance, recording customers while they are filling out the journey map allows to better understand them. Multiple types of multimedia can be used such as audio, video, photos, and sketches.

[00113] Figure 3 shows an example journey map model 300. The model 300 includes an example XML data structure to store journey maps using concepts (journey, customer, touchpoint, experience, channel, stage, lens) arranged in a hierarchical structure. The structure can be used as an import log for process mining. The platform 100 can generate the journey map model 300.

[00114] The example delineates a hypothetical situation in the retail industry. Alice (a customer) called to complain. Therefore, the top management decided to map her journey based on historical data. They observed three events: (1) Alice asked for a quotation, (2) she received the quotation, and (3) the call she made to complain. Example code may be:

```
<cjm>
  <string key="    " value="Understand why sales processes take so long"/>
  <journey>
    <string key="    " value="actual"/>
    <customer>
      <string key="    " value="Alice"/>
    </customer>
    <touchpoint>
      <date key="         " value="2016-05-07T10:26:10"/>
      <string key="    " value="Asking for a quotation"/>
    </touchpoint>
    <touchpoint>
      <date key="         " value="2016-06-08T10:18:03"/>
      <string key="    " value="Receiving the quotation"/>
      <lens>
        <string key="    " value="observation"/>
        <string key="       " value="The duration is too long"/>
      </lens>
    </touchpoint>
    <touchpoint>
      <date key="         " value="2016-06-08T11:52:12"/>
      <string key="    " value="Calling the helpdesk"/>
      <string key="       " value="Complaining about delay"/>
      <experience>
       <string key="       " value="Frustration"/>
       <int key="     " value="-2"/>
      </experience>
    </touchpoint>
  </journey>
</cjm>
```

[00115] In order to map the journey, the model 300 has a root element “journey map”. The model 300 allows for the description of a “goal” described within an element “string”. This notation refers to attribute coding. Going down the tree or hierarchical data structure, an element “journey” is a child node to the root element. This journey map model 300 contains only one journey, but the platform 100 can allow for more than one (denoted by * in Figure 3). The model 300 has the “customer” attribute, which can describe Alice. Different levels of detail are provided for each touchpoint that illustrates the flexibility of the model 300. For instance, the first one indicates only the name of the touchpoint and the time, while the second one has a lens “observation”. Finally, the model 300 represents the “experience” using two attributes: emotion, and scale.

[00116] The model 300 can consider front-stage activities; i.e., the ones visible to the customer. The model 300 can consider the path to be potentially cyclic and non-linear (e.g. “Touchpoint.”, excluding gateways (e.g., “XOR”). Customers’ paths can only be influenced, but not controlled. Moreover, a journey represents a path performed by a single customer. If that was not the case (i.e., gateways allow a single path to represent multiple alternative journeys), the multiplication of touchpoints and the freedom that customers have to navigate through them in their preferred orders may lead to a meaningless map showing that anything can happen in any order (referred to as a “flower model” in process mining). These decisions can reduce the complexity of the model by emphasizing the main goal of journey maps.

[00117] Platform 100 can integrate journey map models with process mining and provide a set of tools that support multiple ways to discover, monitor, and improve processes based on real event logs. Platform 100 can provide a link between process models (e.g., BPMN) and the "reality" captured in event logs. We distinguish a "de jure" model from a "de facto" model. The former is normative, as it intends to steer or control the reality. In contrast, de facto models aim to delineate reality. Process mining provides a logic code to analyze or join both worlds. For instance, one can discover a "de facto" model from the event logs. In turn, one can compare this former model with a "de jure" one.

[00118] Figure 4 shows the mapping 400 of the XES concepts to a model generated by platform 100. Once the mapping 400 is done, the data can be analyzed with process mining. For example, a company can use event logs to answer the following question: does the duration between asking the quotation and receiving it affect the customer's experience?

[00119] Figure 5 shows the process mining framework with a journey map model extension 500 that can be stored by platform 100. The top portion shows the process mining framework. The model can be implemented using XES for process mining to integrate journey map with the process mining framework. Platform 100 can combine data on top of models in process mining to provide a basis to tackle customer journeys using a rigorous approach. Respectively, the expected and actual journey maps correspond to the "de jure" and "de facto" process models.

[00120] Figure 6 shows the application of process mining activities 600, 602 on journey map data and results 604. This can involve data pre-processing 600 by platform 100 to generate synthetic journeys using journey map models. The platform 100 can map the journeys to XES. This can involve process mining 602 by platform 100 to create a set of journeys that result in a neutral or positive experience and a set of journeys that result in a negative experience. This can involve hiding touchpoints. The platform 100 generates visualizations or visual elements of the journeys for display at interface application 130.

[00121] The platform 100 can generate models that are easily exploitable by data analytics tools, is extensible to fit a domain-specific application, and it is not tool-dependent. By bringing process mining techniques and journey maps closer together, platform 100 can close the gap between actual and expected journey maps and we shed light on a potential new area of research, which requires further investigations with real-life collections of journey maps.

[00122] The platform 100 uses process mining that integrates with journey map specificities and cluster journeys and their representatives, to predict the next customer's touchpoint, and to navigate among the journeys.

[00123] The platform 100 can use neural networks 112 for natural language processing and sentiment analysis of text data from journey map data to generate tools for process mining.

[00124] The platform 100 can model and used a dynamic Data Warehouse (e.g. at data storage 110) for storing the event traces in order to facilitate fast access and storage of historical (processed) data. The platform 100 can implement a process for clustering the traces using hierarchical clustering. An example trace can be a customer journey map (or journey map). A journey map can indicate or represent different journeys. The clustering can be based on the Levenshtein distance, for example. This has not been used with events that are characterized by several touch-points. **Note to inventors: please expand on this.** The platform 100 can implement a process for selecting specific attributes/touchpoints and guiding the analysis of focused areas of the hierarchical clustering tree through the notion of "goals". **Note to inventors: please expand on this.** The platform 100 can generate visualizations composed of different views that are updated synchronously to allow end-user to navigate within traces. These views fulfill distinct objective: 1) The first one allows to understand the pattern of activities, 2) the second one offers a holistic view of the hierarchical cluster, and 3) the last one is an area to display descriptive statistics.

[00125] The platform 100 can configure interface application 130 to provide a trace interface that uses hierarchical clustering and indexes for interactive navigation through numerous traces stored as event log formats, for example. The interactive navigation by interface application 130 enables exploration of the underlying traces of the whole set of data available to platform 100 or driven by user goals in order to examine events and patterns in specific areas of interest.

[00126] Traces can provide an understanding of the quality of customer experience at an end-to-end level. The ever growing amount of services offered to users for consumption has made the ability to understand their behavior very important. Similarly important is the knowledge extracted by the increasing number of ways organizations interact with their customers; e.g., a customer might visit a physical store, purchase a product online, and provide feedback on social media which can be collected as a holistic data set for use by platform 100. Traces allow for better understanding of a customer's end-to-end experience when using a service by mapping any

interactions with the company (called touchpoints) on a map that ultimately contributes to better understanding and serving customer needs.

[00127] A model for traces can be integrated with process mining techniques. The platform 100 can configure interface application 130 to explore numerous customer journeys at the same time. The platform 100 can take event logs as input, without using any a-priori information and generate visual elements for journeys. Visualizing event logs on traces enables focus on personal customer activities (e.g., by incorporating customer emotions), rather than the "internally-focus problem-solving approach" and traces can incorporate customer journeys that are deemed exceptional behaviors, rather than removing them to increase model readability.

[00128] Representing many customer journeys in an intelligible can be a challenge. There are limits on the number of journeys to be compared (e.g. less than ten), making the overall process relatively straightforward. However, we argue that companies in the service industries tend to deal with hundreds or thousands of journeys. To overcome this challenge and identify different areas of interest, a hierarchical clustering process is employed to segment the original data. The hierarchical nature allows for a top-down navigation of automatically generated groups of similar journeys. Once the clusters are formed, platform 100 and interface application 130 are able to leverage the contextual information that comes along a typical customer journey such as the customers' characteristics, or the emotions. The platform 100 can use indexes in order to explain why the different clusters were generated. The platform 100 can let the users define their own exploration goals, making it the first goal-oriented tool that allows analysts to set a-priori goals to guide their journey exploration.

[00129] The platform 100 can configure interface application 130 with visual representations to show how numerous event logs can be displayed onto traces and can let users navigate into these journeys.

[00130] The platform 100 can let users upload and explore their own dataset using a trace or journey map layout. To limit the number of journeys displayed on the interface application 130 and allow for an intuitive exploration, the platform 100 uses hierarchical clustering.

[00131] Figure 7 shows an example hierarchical clustering tree 700 that can be generated by the platform 100. Each letter represents an activity of an event trace and the tree is built bottom up by merging the activities that are most similar at each iteration. By default, the platform 100 uses a hierarchical process based on a proximity measure that takes into account the order of

activities and is a variant of the Jaccard similarity based on shingles. In the clustering tree 700 (or dendrogram), the journeys seen in the event logs are at the leaf level and as they get merged they form "representative" journeys or event traces. That is, a journey can refer to an event trace. A representative journey is a single pattern of activities whose purpose is to summarize the patterns contained in a cluster. Because the representative journeys at the top of the tree summarize many – potentially distant – journeys, it will tend to show only few activities shared by many. In contrast, the representative journeys closer to the leafs will show more details. The first layers show general patterns and hide less important activities – like a world map would omit small cities. However, interface application 130 allows a drill-down into 'countries' of interest (i.e., pattern of activities), which would redirect to new traces where the previously hidden 'cities' (i.e., omitted activities) will be shown on interface application 130.

[00132] Figure 8 is an example screenshot visible at interface application 130 and points out three views: journey map view 802, tree view 804, and textual view 806 that are available to navigate the clusters. Each of these views fulfills specific objectives. First, the journey map view 802 shows journeys that are in the same cluster. This representation allows to easily compare the pattern of activities. Second, the tree view 804 displays the hierarchical structure of the journey clusters – useful in providing a holistic view of the clusters and current location. Third, a box per cluster view 806 provides a convenient means to display indexes named "salient characteristics". In this example, the salient characteristics is the top 5 results of a chi-square test applied on all the contextual information. For instance, if at a global level (the entire dataset) the number of women is equal to the number of men, it might be surprising to find a cluster with large majority of women. Therefore, this information might come up as one of the top 5 salient characteristics.

[00133] Moreover, the user might be interested in specific characteristics occurring during the journey. For this reason, we allow user-defined goals. For instance, one might be interested in journeys that started by the activity "attending class" experienced by young people.

[00134] The data storage device 110 has a data warehouse model for storing event traces. An event trace has attributes that indicate activities over time, for example. An event trace can be a set of touchpoints or a sequence of activities. The data storage device 110 can store event logs that represent a collection of event traces. The event traces can be used by platform 100 to generate journeys that can be actual journeys or representative journeys. The event traces can be based on real-time (e.g. current) and/or historical data.

[00135] The data visualization tool 116 can generate visual elements of an interface 800 for an interactive interface application 130 by generating clusters for the event traces using the clustering engine 118 to group event traces having similar attributes. The clustering engine 118 can cluster similar traces that are close together using distance measures. The clustering engine 118 can generate the hierarchical cluster by computing a distance measure for the attributes of the event traces to compute the similar attributes. The clustering engine 118 can generate the hierarchical cluster by computing a distance measure based on a Levenshtein distance. The clustering engine 118 can generate the hierarchical cluster by computing a distance between clusters of the hierarchical cluster. The clustering engine 118 can generate a set of representative attributes for a set of event traces of the hierarchical cluster. The set of representative attributes computed based on the similar attributes. In some embodiments, the pattern of activities indicates the set of representative attributes for a segment of the hierarchical cluster that corresponds to the set of event traces. The clustering engine 118 can generate the hierarchical cluster based on a proximity measure using an order of activities for the event traces.

[00136] The data visualization tool 116 can generate the visual elements for multiple views 802, 804, 806 for the interface 800. A first view 802 indicates a pattern of activities for the set of representative attributes. The first view 802 can indicate representative event traces as an ordered set of activities (home, recreation, shopping, work) over time. A pattern or representative event trace entails many event traces and indicates representative attributes for the activities it represents.

[00137] A second view 804 can indicate the hierarchical cluster and the pattern of activities within the hierarchical cluster. The second view can highlight the location of the pattern of activities in the visual representation of the hierarchical cluster. In some embodiments, the second view indicates the segment of the hierarchical cluster for the pattern of activities indicated in the first view. The clustering engine 118 can generate the hierarchical cluster based on layers and a subset of the layers can be shown in the view 804. A layer can correspond to a number of event traces that will be grouped based on the similar attributes. The layers correspond to a height of the hierarchical cluster. In some embodiments, the clustering engine 118 generates the set of representative attributes based on the similar attributes and a frequent sequences mining process.

[00138] A third view 806 indicates descriptors for the hierarchical cluster. The third view 806 corresponds to the level of granularity of the first and second view 802, 804 and the descriptors

indicate why the subset of the cluster is interesting, for example. The descriptors can indicate distinct traces, total traces and salient characteristics. In some embodiments, the processor computes salient characteristics for the descriptors to indicate indexes based on a chi-square test.

[00139] The multiple views 802, 804, 806 have selectable portions that trigger the data visualization tool 116 to generate updates for the visual elements. The platform 100 controls rendering of the interface application 130 to display the multiple views 802, 804, 806 synchronously. Responsive to a selection of a selectable portion of the plurality of selectable portions, the platform 100 controls rendering of the interface application 130 to update the multiple views 802, 804, 806 synchronously based on the selected portion to navigate the hierarchical cluster. For example, the selectable portions can select attributes (e.g. activities) to filter the views 802, 804, 806 (e.g. event trace contains selected activity X). In some embodiments, the selection portion is within a view of the multiple views to trigger an update to the other views of the multiple views. For example, a selection in view 804 can trigger an update to the other views 802, 804. In some embodiments, the data visualization tool 116 updates the multiple views synchronously based on the selected portion to navigate the hierarchical cluster at different levels of abstraction or granularity.

[00140] In some embodiments, the interface application 130 can show a view indicating actual event traces and a view showing representative event traces. The representative event traces being less than the actual event traces to provide an abstracted representation of the event traces. It may be difficult to see trends and insights from actual events traces if there is too much raw data shown at the interface application 130.

[00141] In some embodiments, the platform 100 receives a navigation goal at the interface application 130. The goal indicates an attribute, for example. Responsive to the navigation goal, the platform 100 controls rendering of the interface application 130 at the device to update the multiple views synchronously to indicate data based on the navigation goal.

[00142] Figure 9 shows a screenshot 900 of interface application 130 for goal setting impact on a tree. The top part displays the settings, while the bottom part shows that some branches of the tree are interesting with regards to the goal (red “hot” area at the top). Hence, the application 130 allows navigation without using any a-priori information, but also setting navigation goals, and guidance by the resulting colors.

[00143] Finally, when moving from one view to the others, the three views are updated on interface application 130 synchronously, allowing a smooth exploration amongst journeys.

[00144] The platform 100 can involve main elements: 1) a web interface; 2) the XES-parser; 3) Hcluster; and 4) a data warehouse. The elements include parameters.

[00145] Web interface. For example, the web interface leverages bootstrap, jquery and d3js to provide a user-friendly interface to upload and navigate journeys. Both the journey map view and the tree view are implemented in d3js. There may be other implementations.

[00146] XES Parser. The platform 100 works with event logs. More specifically, platform 100 can leverage XES (eXtensible Event Stream) within process mining. The XES Parser can be a Java implementation that encapsulates the OpenXES library to parse XES file.

[00147] Hcluster. Hcluster is an implementation containing the three steps for example. Figure 10 illustrates example steps 1000. A step is the hierarchical clustering implemented using Scipy. Two parameters can be provided as inputs: the distance measure between event sequences and the methods for calculating the distance between clusters. They can both be chosen by the user when uploading a dataset. Figure 11 illustrates a preview of the parameters 1010 when uploading a dataset. It seems that using shingles as a distance metric provides a more intuitive way to navigate through journeys. Once the hierarchical cluster is formed, the next step consists of cutting the clustering to form layers. A layer is a set of predetermined number of journeys that will be grouped together and will ultimately appear on the same trace. To achieve this, the platform 100 can recursively cuts the dendrogram returned by Scipy. A small number of journeys will lead to a simple journey map that is easy to visualize, but also works for more complex tree structures (i.e., trees with larger height). Another step consists of finding the representative journey using a frequent sequences mining process.

[00148] Data Warehouse. Each dataset is saved by data storage 110 in its own database schema designed as a star schema. Data storage can store a full schema or all the information required to use the application (e.g., clusters, journeys, events) as well as some precomputations. For instance, the platform 100 can count the number of occurrences for each characteristic at each cluster, so the goals and the salient characteristics can be retrieved quickly.

[00149] Altogether, the parameters visible in Figure 11 allow users to explore the exact same dataset from different perspectives.

[00150] The platform 100 can process many journeys for display onto journey map or traces in an intelligible and efficient manner. The platform 100 integrates process mining taskforce (i.e., XES) and process mining activity (i.e., discovery).

[00151] The platform 100 can cluster events (traces) for the discovery of the number of segments (clusters) that exist. The platform 100 can implement a process that automatically finds the best representatives. This technique is based on genetic processes. Except for the timing and the title of each activity in a trace of events, the platform 100 can take into account any contextual information of the activities, for example looks into any text the users wrote or any emotion they showed.

[00152] In some embodiments, clustering engine 118 can compute clusters of the event traces based on a genetic process function. In some embodiments, for each event trace, the attributes are represented as a two-dimensional tuple of activities over time and contextual data and the genetic process function can evaluate the two-dimensional tuples of the event traces.

[00153] Each cluster can correspond to a set of event traces and a representative event trace based on representative attributes of the set of event traces. The genetic process function maps each event trace to a cluster. The data visualization tool 116 can generate the visual elements for the interface application 130 and the visual elements can indicate the clusters and, for each cluster, the representative event trace based on the representative attributes. The representative event trace summarizes the set of event traces corresponding to the cluster. The data visualization tool 116 can control rendering of the interface application 130 at the device to display the visual elements and a plurality of selectable portions. Responsive to a selection of a selectable portion of the plurality of selectable portions, data visualization tool 116 can control rendering of the interface application 130 to navigate the clusters. The clustering engine 118 can determine segments for the event traces and then evaluate them to determine the clusters, for example. The clustering engine 118 can group the event traces into X groups and then identify a representative event trace for each group. The representative event trace does not have to be an actual event trace or representative of actual raw data but can summarize the group or cluster of event traces it represents. In some embodiments, the platform 100 can use neural networks 112 for natural language processing of text data that can form part of event traces to determine a sentiment score or sentiment indicator. This sentiment score or sentiment indicator can be an attribute for the event trace, for example. The genetic approach goes through different

combinations of clusters and evaluates the clusters to determine the best segments of event traces.

[00154] In some embodiments, clustering engine 118 can implement the genetic process function by: evaluating initial set representative event traces to generate elite set representative event traces, generating additional initial representative event traces using a transformation process, evaluating the additional initial representative event traces to generate additional elite set representative event traces, continuing the generating and the evaluating until stopping criterion is met, using resulting elite representative event traces as the representative event traces for the clusters.

[00155] In some embodiments, clustering engine 118 computes, for each cluster, the representative event trace as a pattern of activities to summarize the patterns contained within the set of event traces of the cluster. In some embodiments, clustering engine 118 implements the genetic process function by pre-processing the event traces, determining an initial population that corresponds to a number of the clusters, evaluating initial representative event traces, evaluating stop criteria, implementing a genetic operation to generate additional initial representative event traces, outputting a map abstraction for use in generating additional visual elements. In some embodiments, clustering engine 118 is configured to evaluate the initial representative event traces based on quality criteria for fitness, a number of representatives, a contextual distance and an average quality. In some embodiments, the clustering engine 118 is configured to weigh the quality criteria to generate an overall quality and average the overall quality. In some embodiments, the clustering engine 118 is configured to evaluate the initial representative event traces based on internal and external evaluation metrics.

[00156] Summarizing numerous customer trajectories is challenging. The platform 100 uses clustering to summarize sequences of categorical data. Using synthetic datasets simulating customer journeys, as well as a real dataset, the platform 100 is flexible and returns customer journeys of better quality.

[00157] A customer experience can be defined as a customer's journey with a firm over time across multiple interactions called touchpoints. A journey can be as simple as a single activity (e.g., 'looking at a product'), but can also involve complex interactions through various channels, multiple devices, and from several locations. Customer interactions are increasing, services are becoming more complex, and customers can be unpredictable. Moreover, there is an increasing number of customer-to-customer interactions, which are drivers of variability and unpredictability.

Overall, offering a great customer experience is a challenging task, partly due to the lack of conceptual clarity. Furthermore, a lack of a unified understanding of customers can be an obstacle to mobilizing employees around integrated touchpoints, journeys, and consistent experiences.

[00158] Concretely, a challenge faced by many practitioners is to make sense of the–potentially infinite–combination of activities that exist in order to consume a service. A trace is a conceptual tool used to better understand customers' trajectories when they are consuming a service. It depicts typical journeys that will be experienced by the customers across several touchpoints. Multiple types of journeys can be included on a trace. When a trace is used as a design thinking tool by internal stakeholders to discuss, challenge, or innovate the way a service is offered to customers, expected journeys are used. This implies a dyadic relation in which it cannot be assumed that the experience is perceived by customers as intended by the internal stakeholders. A trace can also be used to understand how the service is really perceived and experienced by customers, either by directly eliciting their opinion through surveys and special tools, or by leveraging evidences stemming from data. In this case, it can be referred to it as the actual journey.

[00159] The platform 100 can transform a set of actual journeys into a few representative journeys with the aim of displaying them on an interface application 130 as visual elements. Figure 12 is an example screenshot 1200 of interface application 130. This task is illustrated in the four parts of Figure 12.

[00160] For the sake of the example, the screenshot 1200 defines five typical touchpoints 1202 with which a customer might interact when buying a product Then, screenshot 1200 can depict five journeys by event logs 1204 composed of three to four touchpoints. We consider these journeys to be actual because they represent what the customers have really experienced. Typically, the data can be extracted from software data sources 160. Moreover, the data can be formatted according to the XES format for process mining. The important point is to be able to group the touchpoints from the same journey together (using a journey identifier). XES format can store customer journeys. Next, the screenshot 1200 can display the five actual journeys 1206 from the event logs on a journey map or trace. Although five journeys can still be considered as a limited number of journeys that can display behaviors. This can reduce the complexity and looks at these journeys at a higher level of abstraction. Specifically, a representative journey 1208 is a single pattern of activities whose purpose is to summarize similar actual journeys, which is a

technique useful for characterizing typical trajectories. This shows how the usage of two representative journeys 1208 helps in reducing a journey map's or trace's complexity. As it can be seen, each actual journey is assigned to its closest representative journey, e.g., the darker representative journey summarizes the actual journeys 2 and 5. In this regard, finding representative journeys can be considered as a clustering task by platform 100.

[00161] Finding representative journeys involves platform 100 choosing the right number of journeys and finding a sequence of activities that summarizes the actual journeys well. Moreover, the results can depend on the underlying motivation for building the journey map in the first place. For instance, if the goal is to have a general overview of the actual journeys, there may be a different journey map than if the goal were to investigate peculiar arrangements of touchpoints. Finally, one may wish to find groups of similar customers solely based on their trajectories (i.e., the sequence of touchpoints), or based on their characteristics (e.g., age and region), or most probably, on a continuum between these two extremes.

[00162] Given the challenges surrounding the desire to increase the understanding of customer behavior, the platform 100 can be configured for clarifying the customer journey discovery activity, proposing ground truth datasets, which are particularly suited for evaluating this activity, and introducing a process to discover representative journeys. Using the proposed datasets and existing cluster analysis techniques, platform 100 can generate traces automatically in a flexible nature using a real dataset and illustrate the results at interface application 130.

[00163] The customer journey discovery activity can be described with the following definition: given a set of actual journeys, find a reasonable amount of representative journeys that well summarizes the data. The following table provides a summary of notation:

t	touchpoint
$\|T$	finite set of all touchpoints
S	nonempty and ordered set of touchpoints
C	set of contextual data
c	name of the contextual data
$=c$	nominal value of the contextual data
J_a	an actual journey observed from a customer
$J_{\mathcal{A}}$	event logs, set of all the actual journeys
J_r	a representative journey summarizing a subset of $J_{\mathcal{A}}$
$J_{\mathcal{R}}$	customer journey map, a set of J_r summarizing $J_{\mathcal{A}}$
$k_{\mathcal{R}}$	number of J_r in $J_{\mathcal{R}}$; i.e., $\|J_{\mathcal{R}}\|$
$\sigma_{\mathcal{A}}$	set of all the S observed in event logs
$\sigma_{\mathcal{R}}$	set of all the S contained in $J_{\mathcal{R}}$
$\kappa_{\mathcal{A}}$	set of all the C observed in event logs
$\kappa_{\mathcal{R}}$	set of all the C contained in $J_{\mathcal{R}}$

Table 1: Summary of the notation.

[00164] A touchpoint is an interaction between a customer and a company's products or services. 'Sharing on social network' or 'ordering a product on the website' – two activities visible in Figure 12 – are two typical examples of touchpoints in a retail context. 'Visiting the doctor', or 'taking medication' are other possible examples in a healthcare context. Let t be a touchpoint, and let T be the finite set of all touchpoints.

[00165] Ordered Set of Touchpoints involves one or more touchpoints. When consuming a service, a customer will have a trajectory composed of one or more touchpoints. Let $S = t_1, t_2, ...; t \in T$ be a set of ordered touchpoints. For instance, at 1204 the first journey starts with 'visiting the shop', followed by 'testing the product', and finally 'sharing on social network'.

[00166] Contextual data is the data that holds contextual information relevant during the service delivery. Typically, customers' demographics or the level of satisfaction of the customers are examples of contextual data. For instance, we might discover that certain trajectories are more popular amongst younger customers. In this case, it would constitute important information for the marketing team in order to better anticipate and communicate with this segment of customers. The intuition is that discovering trajectories that can be described by particular contextual data (e.g., the trajectories of younger customers from the 'region A') provide more insights than finding representatives that are uniformly distributed amongst all customers' profile. Hence, adding the contextual data in the journey allows to consider it when trying to find the best representative journey. Let C be a set of all available contextual data. Let the name of the contextual data be $c \in C$ and let $\#c$ be its nominal value.

[00167] An actual journey J_a can be a tuple (S, C); i.e., an actual journey is a nonempty and ordered set of touchpoints and a set of contextual data. For instance, J_a = {('Visiting the shop', 'Testing the product', 'Sharing on social network'), {subscription-type:standard}, {age-range:young}} is a journey with three touchpoints from a young customer having a standard subscription.

[00168] J_A can refer to an event log, the set of all actual journeys observed from customers.

[00169] A representative journey summarizes a subset of J_A. Let this subset be J_A. Let a representative journey J_r be a triplet $(S, C, J_A \subseteq J_A)$; i.e., a representative journey is a nonempty and ordered set of touchpoints, and a set of contextual data, which summarizes a subset of J_A. At 1208 there is shown two representative journeys summarizing five actual journeys. As can be

seen, the lighter representative journey ⟨*BDAE*⟩ represents the actual journeys ⟨*BDA*⟩, ⟨*BDAE*⟩, and ⟨*BDE*⟩.

[00170] A customer journey map is a conceptual tool used to provide an overview of the customer experience through the use of representative journeys. It contains one or several representative journeys. Let a customer journey map JR be the superset of all the J_r. Let k_R denotes the number of journeys on a map (i.e., $|J_R|$).

[00171] Finally, for readability purposes, let $\sigma_A = \{S : S \in J_a, J_a \in J_A\}$ and $\sigma_R = \{S : S \in J_r, J_r \in J_R\}$; i.e., σ_A is the set of all the sequences of activities observed from customers, while σ_R is the set of all the sequences of activities displayed on a customer journey map. Similarly for contextual data, let $\kappa_A = \{C : C \in J_a, J_a \in J_A\}$ and $\kappa_R = \{C : C \in J_r, J_r \in J_R\}$.

[00172] The customer journey discovery activity can be defined as a function that maps all members of J_A to a member of J_R; i.e., that maps all the actual journeys to representative journeys ultimately displayed on a journey map. Note a journey J_a is assigned to a single representative journey in J_R.

[00173] Discovering J_R from J_A is an unsupervised clustering task implemented by clustering engine. The platform 100 determines a parameter for the number of k_R. When the goal is to have a general overview about J_A, it seems reasonable to have k_R in a range from two to six journeys so the journey map is readable but this can vary. However, discovering few dozens J_r might also be a relevant choice if the goal is to catch complex and less generic patterns. Once k_R has been found, the another interesting challenge lies in the fact that J_a is two-dimensional. Indeed, both *S* and *C* can be taken in consideration when clustering the journeys. Finally, the sequence $S \in Jr$ that best summarizes its assigned actual journeys needs to be found. It might be the case that an ideal representative journeys was never observed but well summarizes the actual journeys.

[00174] The platform 100 uses a hierarchical structure so that the first layers show only the most important activities, abstracting from less important ones. The platform 100 can explore without any a priori information, but it is also possible to set goals based on journey characteristics (e.g., age and gender). Once the goals are set, areas of interest (i.e., the ones fulfilling the goal) are highlighted at interface application 130. In platform 100, the dendrogram is built based on the edit distance between journeys. The platform 100 can also leverage the contextual information when finding representatives.

[00175] Figure 13 is a diagram 1300 that shows that the process mining framework as well as the input data are relevant in a customer journey analytics context. Similar to the process discovery activity in process mining, platform 100 focuses on the discovery of a model from event logs. Figure 14 shows that these business process model (BPM) 1402 and journey maps 1404 convey different information. A journey map depicts journeys as experienced by customers while a BPM shows the available combination of activities using advanced constructs such as XOR or parallel gateways. Finally, the information that will be leveraged for a journey map differs from that used for a BPM. For instance, the satisfaction, the emotion, and various customer characteristics are central pieces of information that should be leveraged when discovering a journey map. Such information might sporadically be used to enhance a BPM, but usually only once the model has been discovered.

[00176] The platform 100 can summarize a set of events using representative sequences summarizing a set of sequences. There are different ways to select a representative. The 'frequency', where the most frequent event is used as the representative. The 'neighborhood density', which consist of counting the number of sequences within the neighborhood of each candidate sequence. The most representative is the one having the largest number of sequences in a defined neighborhood diameter. The 'mean state frequency': the transversal frequencies of the successive states is used to find a representative using the following equation:

$$MSF(s) = \frac{1}{\ell} \sum_{i=1}^{\ell} fsi \qquad (1)$$

where

$$s = s_1 s_2 \ldots s_l : \text{Sequence of length } \ell$$

$$f_{s1}, f_{s2}, \ldots f_{sl} : \text{Frequencies of the states at time } t_\ell$$

[00177] The sum of the state frequencies divided by the sequence length becomes the mean state frequency. Its values is bounded by 0 and 1 where 1 describes a situation where there is only one single distinct sequence [9]. The sequence with the highest score is the representative. The 'centrality': the representative – or medoid – can be found using the centrality. Being the most central object, the representative is the sequence with the minimal sum of distances to all other sequences. Finally, the 'sequence likelihood': the sequence likelihood of a sequence derived from the first-order Markov model can also be used to determine the representative.

[00178] The platform 100 can use genetic approach that uses contextual data and can be an extension to the summarization of categorical sequences.

[00179] The platform 100 can use genetic approaches to discover business process models from event logs. The platform 100 tailors it towards journey map by introducing specific evaluation metrics suited to measure the quality of a journey map and its representative journeys given a set of actual journeys. Figure 15 depicts the main phases 1500: (1) a preprocessing phase, (2) generating the initial population, (3) assigning each actual journey to its closest representative, (4) evaluating the quality of the journey maps, (5) the stopping criterion evaluation, and (6) creating new journey maps by applying some genetic operations.

[00180] As an introduction, Figure 16 provides some intuitions on the genetic process. Typically, a set of actual journeys, σ_A, is provided to the algorithm. First, during *Generation 1* 1602, the platform 100 can build p number of J_R, where p is the population size; i.e., the number of J_Rs that will be evaluated after each generation. In our example $p = 3$. Each J_R is evaluated and we keep the e best J_Rs, called the *elite set* while we discard the other $(p - e)$ J_Rs. Then platform 100 can move to *generation* 2, 1604 where it can keep an untouched version of the e number of J_Rs in the elite set. In Figure 16, if $e = 1$, platform 100 can keep the best J_R from generation 1, i.e., J_{R2} is intact in generation 2. The platform 100 can then apply some transformation to generate $(p - e)$ new J_Rs that will, in turn, be evaluated. The platform 100 can generate J_{R4} and J_{R5}. The platform 100 can recursively transform and evaluate the p number of J_Rs until a stopping criterion is met. Once a stopping criterion is met, we return the best J_R (J_{R7} in Figure 16 at generation 3 1606). The best J_R can be interpreted as the best set of representative journeys, J_r, representing a set of journeys from J_A that have been found given certain evaluation criterion. The platform 100 can generate the initial population, apply various types of operations on each J_R to transform them, and evaluate each one of them given a set of journeys in J_A.

[00181] The platform 100 can make the assumption that σ_R will be close to the frequent patterns observed in σ_A.

[00182] Let $\sigma_{A\ell}$ be the set of all actual journeys of length ℓ and let $Top_\ell = \{Top_\ell \in \sigma_{A\ell}, |Top_\ell| = \min(|\sigma_{A\ell}|, in)\}$ be the i_n most occurring pattern of length ℓ. Finally, let $Top = \{Top \in Top[i_s, i_e]\}$, i.e., Top is the superset of all the most occurring patterns of lengths i_s to i_e. Figure 17 depicts Top 1700 from the running example given in Figure 16.

```
Input  : σ_A, actual journeys
         p, population size
         e, elite size
         setQuality, set of quality criteria
Output : J_R, best customer journey map

1  Top ← findTopPattern(σ_A)

2  setCjm ← Array()
3  for i ← 0 to p do
4  |  AJ ← pickRandom(Top)
5  |  setCjm.add(J_R)
6  end
7  nextPopulation(setCjm)
8  Function nextPopulation (setCjm)
9  |  foreach J_R in setCjm do
10 |  |  assignToClosestRepresentative(J_R, σ_A)
11 |  |  evaluateQuality(J_R, setQuality)
12 |  end
13 |  if stopCriterionMet(setCjm) then
14 |  |  return topRepresentative(setCjm, 1)
15 |  else
16 |  |  elite ← topRepresentative(setCjm, e)
   |  |  setNextCjm ← elite
17 |  |  for i ← 0 to p do
18 |  |  |  AJ ← pickRandom(elite)
19 |  |  |  setNextCjm.add(J_R)
20 |  |  end
21 |  |  nextPopulation(setNextCjm)
22 |  end
```

[00183] *Top* is used later to form the initial population of J_R, and to add a random journey to J_R. Using *Top* the platform 100 can avoid generating journeys by picking a random number of touchpoints from *T*. Using *Top* can reduce the execution time by two to get an output J_R of the same average quality. For an example, platform 100 can fix i_n to 10, i_s to 2, and i_e to 12.

[00184] To build the initial population platform 100 can generate *p* number of J_R on which we add one sequence randomly picked from Top (defined in Sect. 4.1). This is visible in the genetic process shown, lines 2 to 5. As can be seen in Figure 16, $k_R = |J_{R1}| = |J_{R2}| = |J_{R3}| = 1$ because platform 100 can only add one sequence per J_R.

[00185] To evaluate how good J_R describes J_A, platform 100 can decompose the problem. Indeed, the quality of J_R can be based solely upon its representative journeys. And, the quality of

a representative journey can be measured when knowing which actual journeys it represents. Hence, a first step toward evaluating the quality of J_R is to assign each journey $J_a \square J_A$ to its closest journey in $J_r \square J_R$. This is illustrated in Figure 12 at 1208 where each journey J_a is assigned to a journey J_r.

[00186] To characterize the closeness between J_a and J_r, platform 100 can use the Levensthein distance. It is a metric to measure the distance between sequences. The Levensthein distance counts the number of edit operations that are necessary to transform one sequence into another one. There are three types of operations: deletions, insertions, and reversals (or substitutions). For instance, the distance between ⟨ABC⟩ and ⟨ACCE⟩ is 2 since one substitution and one insertion are required to match them.

[00187] The platform 100 can define the closest representative as the one having the smallest Levensthein distance with the actual journeys. Note that if a tie occurs between multiple best representatives, platform 100 can assign the J_a to the J_r having the smallest amount of actual journeys already assign to it. Moreover, when $k_R = 1$ all the J_a will be assigned to the same J_r. This be the case after creating the initial population because only one J_r is added to J_R.

[00188] Once each actual journey has been assigned to its closest representative, platform 100 can evaluate J_R using the criteria described in the next section.

[00189] The platform 100 can use different evaluation criteria to determine the quality of each J_R, namely, (1) the fitness, (2) the number of representatives, (3) the contextual distance, and (4) the average quality. Using the example in Figure 16, Figure 18 illustrates the evolution 1800 of the evaluation criteria for the best J_R.

[00190] The fitness measures the distance between each sequence of activities σ_A and its closest representative σ_R using the Levenshtein distance.

$$Fitness(\sigma_{\mathcal{A}}, \sigma_R) = 1 - \frac{\sum_{i=1}^{|\sigma_{\mathcal{A}}|} min_{j=1}^{\sigma_R} \left(Levensthein(\sigma_{\mathcal{A}_i}, \sigma_{R_j})\right)}{\sum_{i=1}^{|\sigma_{\mathcal{A}}|} Length(\sigma_{\mathcal{A}_i})} \quad (2)$$

where

$\sigma_{\mathcal{A}_i}$	: i^{th} actual sequence observed in event logs
σ_{R_j}	: j^{th} representative contained in J_R
$Levensthein(x_1, x_2)$	: Levensthein distance between two sequences
$Length(x)$	: Length of the sequence of activity x

[00191] A fitness of 1 means that the representative journey perfectly catches the behavior of the actual journeys assigned to it. In contrast, a fitness close to 0 implies that many edit operations are necessary to match the sequences.

[00192] Another evaluation criteria is the number of representatives. An event log can contain several hundreds or thousands of unique actual journeys. Hence, if platform 100 maximizes the fitness without trying to keep a low k_R, the journey map can become unreadable because too many representative journeys may be displayed on it. In other words, J_R overfits. Hence, the goal is to find a k_R that offers a good compromise between underfitting and overfitting. Finding the optimal number of clusters is a recurrent challenge when clustering data. The platform 100 can involve integrating different ways of determining the ideal number of clusters, such as the Bayesian information criterion, the Calinski-Harabasz index, or the Silhouette technique. The idea is to evaluate a range of solutions (e.g., from 2 to 10 journeys) and to keep the best solution. Let k_h be the ideal number of clusters returned by one of the three techniques mentioned above. By integrating k_h into the evaluation, platform 100 can guide the solution toward a k_R that is statistically relevant. To evaluate the quality, platform 100 measure the distance between k_R and k_h. To do this, platform 100 can use the following example distribution function:

$$NumberOfRepresentatives(k_R, k_h, x_0) = \frac{1}{1 + (\frac{|k_R - k_h|}{x_0})^2} \quad (3)$$

where

k_R : Number of J_r journeys on J_R (i.e., $|J_R|$)
k_h : Optimal number of journeys (e.g., using the Calinski-Harabasz index)
x_0 : x value of the midpoint

[00193] Figure 19 illustrates a graph 1900 of the distribution for kh=6 and x_0=5. For instance, if k_R = 2, it would have a quality of ≈ 0.6. Also note that the parameter x_0 determines where the midpoint of the curve is. Concretely, k_R = 1 or k_R = 11 can result in a quality of 0.5 because the absolute distance from k_h is 5. The platform 100 set x_0 = 5 for experiments.

[00194] The contextual distance allows platform 100 to consider the set of contextual data C when grouping similar journeys. The more distant the set of contextual data is between J_a that are represented by distinct J_r, the better the quality is. To measure the distance, platform 100 can first build a value frequency table which count all the values per representatives (see Figure 20). Then, for each pair of clusters, platform 100 can calculate the cosine similarity, which is defined as:

$$ContextualDistance(v_1, v_2) = \frac{v_1 \cdot v_2}{||v_1|| \cdot ||v_2||} \quad (4)$$

where

v_1 : Value frequency counter for J_{r1}
v_2 : Value frequency counter for J_{r2}

[00195] Finally, the cosine distances are average to get the overall contextual distance. A short overall distance indicates that the contextual data of Ja that are assigned to distinct Jr are similar. In other words, the contextual data does not help in classifying Ja between several Jr.

[00196] The quality criteria can be weighted by platform 100. Then, platform 100 can average the overall quality as follows:

$$AvgQuality(w_f, w_{k_h}, w_c) = \frac{(F * w_f) + (N * w_{k_h}) + (C * w_c)}{w_f + w_{k_h} + w_c} \quad (5)$$

where

F : Fitness quality
N : k_h quality
C : Contextual Distance quality
w_f : Weight for the quality criteria 'fitness'
w_{kh} : Weight for the quality criteria 'number of representatives'
w_c : Weight for the quality criteria 'contextual Distance'

[00197] A weight of 0 skips that criterion. The results can be best if more weight is given to the fitness quality. Typically, weights $w_f = 5$, $w_{kh} = 1$, and $w_c = 1$ lead to the best results both for the synthetic datasets and during our experimentation with a real dataset.

[00198] Once all the J_R from p have been evaluated, platform 100 can rank them by decreasing quality.

[00199] Before creating new J_R, platform 100 can make sure that a stopping criterion is not met. There are three example evaluation criteria the platform 100 can use to stop: (1) The platform 100 could stop after a certain number of generations. (2) platform 100 could stop when a certain number of generations have been created without improving the average quality. (3) platform 100 could stop when a certain quality threshold is reached for one of the evaluation criteria. Because it is difficult to predict the quality level that can be reached, platform 100 might not stop using a threshold. For this reason, platform 100 can use a combination of approaches.

[00200] Once the stopping criteria have been evaluated, there are two outcomes. Either one of the criteria is met and the platform 100 stops, returning the best J_R (algorithm 1, line 14), or, platform 100 generate new candidates by recursively calling the function nextPopulation of the algorithm 1 (line 21). How platform 100 generates new candidates is described herein.

[00201] Once all the journey maps have been evaluated, platform 100 can rank them by their average quality and copy a fraction (i.e., e) of the best ones. Because platform 100 can keep an untouched version of the e number of $J_R s$, platform 100 can make sure that the overall quality will only increase or stay steady. Indeed, as can be seen in Figure 18, the average quality (i.e., black line) is never decreasing.

[00202] Then, platform 100 can generate $(p - e)$ new $J_R s$ as follows. The platform 100 can pick one random J_R from elite, and perform one or multiple operations. The platform 100 can define four different operators: (1) add a journey, (2) delete a journey, (3) add a touchpoint, and (4) delete a touchpoint. The way they are applied is determined using algorithm 2. In extreme cases, platform 100 might execute each of these four operators three times. In any case, platform 100 perform at least one operation (enforced by line 2 in Alg. 2). Typically, in Figure 16, J_{R4} have been produced by taking J_{R2} and applying two operations: (1) removing the touchpoint 'D', and (2) adding a journey picked from *Top*.

[00203] A sequence is randomly picked from *Top* and added to J_R. For instance, in Figure 16, J_{R7} has been produced by taking J_{R5} from *Elite* and adding a journey from *Top*.

[00204] A random journey is removed from J_R. Nothing happens if J_R contains only one journey.

[00205] A touchpoint from T is added to one of the journeys from J_R at a random position. For instance, in Fig. 5, J_{R6} has been produced by adding the touchpoint 'D' at the second position of J_{R5}.

[00206] A touchpoint is removed from J_R unless removing this touchpoint would result in an empty set of touchpoints. The following shows application of the genetic operations.

```
transformed ← False ;
while transformed == False do
    for i ← 0 to 2 do
        if randFloat(.0, 1.) <= .1 then
            addJourneyFromTop()
            transformed ← True
        end
        if randFloat(.0, 1.) <= .1 then
            if |M| > 1 then
                deleteRandomJourney()
                transformed ← True
            end
        end
        if randFloat(.0, 1.) <= .1 then
            journey ← pickRandomJourney()
            journey.addRandomTouchpoint()
            transformed ← True
        end
        if randFloat(.0, 1.) <= .1 then
            journey ← pickRandomJourney()
            if |journey| > 1 then
                journey.deleteRandomTouchpoint()
                transformed ← True
            end
        end
    end
end
```

[00207] As described in Fig. 15, once new J_Rs have been created, platform 100 can go back to the evaluation phase where the new J_Rs are evaluated until one stopping criterion is met. If such a criterion is met, platform 100 can return the best J_Rs of the last generation (e.g., J_{R7} in Fig. 16), putting an end to the execution.

[00208] In order to evaluate the quality of the approach to return the best set of representative journeys in JR, platform 100 can evaluate the results using a collection of synthetic datasets. The platform 100 can generate the dataset. Then, using this synthetic dataset, platform 100 can evaluate and compare to summarize sequences of categorical data. In this first evaluation, platform 100 does not leverage contextual data to allow for a fair comparison that consider only the sequence of activities. Finally, platform 100 can add contextual data to show the gain in information the platform 100 can get from leveraging such data.

[00209] To evaluate the results of the approach, platform 100 can produce several event logs that simulate journeys. The platform 100 can generate the event logs using the ground truth represented by the generative journeys and recover these journeys from a set of actual ones it produces. A generative journey is a known sequence of activities with a known set of

characteristics from which we generate the event logs. Similar to J_a, a generative journey *Jg* is a tuple (*S, C*), containing a nonempty and ordered set of touchpoints, *S*, and a set of contextual data, *C*. Let J_G be a set of k_G generative journeys used to generate a dataset composed of 1,000 actual journeys. Let $\sigma_G = \{S : S \in J_g, J_g \in J_G\}$; i.e., σ_G is the set of all the sequences defined in the generative journeys.

[00210] If platform 100 were to use only these generative journeys to generate 1,000 thousand journeys, platform 100 would obtain only k_G distinct journeys. For instance, if platform 100 uses j_{g1} = {⟨ABC⟩, {age-range:range1}} and j_{g2} =(⟨ABBD⟩, {age-range:range2}) to generate 1,000 journeys equally distributed, we will obtain $j_a = \{j_{g1}^{500}, j_{g2}^{500}\}$. From a business point of view, this would describe an ideal situation where each group of customers behaves in an homogeneous way. A more realistic situation would depict a scenario where each group of customers can be described by a representative sequence of activities, but the actual journeys within the group can deviate from the representative one. To produce more realistic data, platform 100 can inject noise for a fraction of the journeys. For instance, if the noise level is set to 50%, $J_a = J_g$ is true for half of the data. For the other half, algorithm 3 describes how noise is added. Since the noise is added to the two components of the journey, there are two parts: the sequence of touchpoints and the contextual data. Figure 21 illustrates how six journeys 2104 are generated from two generative journeys 2102. If we assume that the noise level is defined to be 50%, three actual journeys in the events logs deviate from the original generative journeys. The following shows operations for adding noise to the sequence of activities (e.g. event traces)

```
initialSequence ← twistedSequence ;

while initialSequence == twistedSequence do
    for i ← 0 to 5 do
        if randFloat(.0, 1.) <= .2 then
            twistedSequence ← removeRandomActivity()
        end
        if randFloat(.0, 1.) <= .2 then
            twistedSequence ← addRandomActivity()
        end
        if randFloat(.0, 1.) <= .2 then
            twistedSequence ← swapTwoActivity()
        end
    end
end

foreach char in characteristics do
    char ← chooseRandomObservedValue() ;
end
```

[00211] The goal of experiments is to retrieve the set of generative journeys, as representatives, from the produced actual journeys. The noisier the dataset is, the more difficult it is to retrieve the generative journeys from the actual journeys. However, the noise is not the only source of complexity. Hence, to ensure complexity-wise diverse datasets, platform 100 can define three additional sources of complexity:

(1) Complex number of touchpoints: The more touchpoints are contained in the generative journeys, the more difficult it is to retrieve them. We define that a journey map possessing 20 or more touchpoints is considered complex.

(2) Complex contextual data: The more similar the contextual data is between generative journeys, the more difficult it will be to distinguish them. Let c be the average contextual distance (defined in Sect. 4.4.3) when there is no noise. We define that if $c \geq 0.25$, we consider the journey as complex.

(3) Complex distribution: If the journeys in the event logs are not well distributed among the generative journeys, it will also be harder to retrieve the original generative journeys. We define that if the standard deviation of the distribution (normalized to 1) is equal to or greater than 0.25, we consider the journey as complex. Typically, if 300 journeys are generated from J_{g1} while 700 journeys are generated from J_{g2}, we would consider the journeys as complex (i.e., the standard deviation is 0.283).

[00212] The binary combination of these three sources of complexity results in eight different datasets, which are described in Table 2. As can be seen, the first row (i.e., #0) describes a configuration with 4 generative journeys (column '*kG*'), a total of 17 touchpoints, a contextual distance of .17, and a standard distribution of .17 (column 'Distribution'). Hence, we do not consider the row #0 as complex for any of the three complexity mentioned above (column 'Complex'). Combining these 8 configurations with 5 levels of noise produces 40 datasets as an example experiment.

#	Complex (1)(2)(3) (t=true)	k_G	Touch-points	Contextual distance*	Distribution (std_dev)
0		4	17	0.17	[.10, .10, .30, .50] (.17)
1	t	5	25	0.00	[.10, .10, .20, .20, .40] (.11)
2	t	4	17	0.33	[.10, .15, .10, .65] (.23)
3	t	3	11	0.00	[.10, .10, .80] (.33)
4	t t	4	20	0.31	[.25, .25, .25, .25] (.00)
5	t t	3	13	0.25	[.08, 0.10, 0.82] (.34)
6	t t	6	27	0.00	[.05, .05, .05, .05, .05, .75] (.26)
7	t t t	5	21	0.25	[.07, .09, .09, .05, .7] (.25)

*average of the cosine distance when there is no noise

[00213] To evaluate and compare the quality of representative journeys, platform 100 can use both external and internal evaluation metrics. On the one hand, the external ones evaluate the results relative to the ground truth, i.e., from the generative journeys. On the other hand, the internal evaluation uses cluster analysis techniques to assess the results. The aim is to account for the fact that the ground truth might not be the optimal solution. Indeed, adding random noise might change the optimal solution.

[00214] An example metric for External Evaluation is Distance in Number of Journeys. This metric measures the distance between the number of generative journeys and the number of representative journeys returned by the algorithm. The platform 100 can use the following metric:

$$NbJourneysDistance(k_G, k_R) = abs(k_G - k_R) \qquad (6)$$

[00215] An example metric for External Evaluation is Jaccard Distance. To evaluate the distance between the sequences of activities from the generative journeys (σ_G) and the discovered representative journeys (σ_R), platform 100 can use the Jaccard Distance where a score of 1 indicates a perfect match between the set of sequences from the generative journeys and the representative ones.

$$JaccardDistance(\sigma_R, \sigma_G) = 1 - \frac{|\sigma_R \cap \sigma_G|}{|\sigma_R \cup \sigma_G|} \qquad (7)$$

[00216] An example metric for Internal Evaluation is Mean distance. The mean distance *i* returns the average distance between the representative sequence *i* and the sequence of actual

journeys that have been assigned to *i*. If the mean distance *i* is 0, then the representative journey *i* perfectly matches the underlying actual journeys.

$$MeanDistanceScore_i = \frac{\sum_{j=1}^{k_i} D(S_i, S_{ij})}{k_i} \quad (8)$$

where

$D(x_1, x_2)$: Levensthein distance between two sequences
k_i : Number of actual journeys assigned to the representative journey i
S_i : Representative sequence i
S_{ij} : Sequence of actual journeys j assigned to i

[00217] An example metric for Internal Evaluation is Coverage. The coverage indicates the proportion of actual journeys that are within the neighborhood n of a representative.

$$Coverage_i = \frac{\sum_{j=1}^{k_i} (D(S_i, S_{ij}) < n)}{k_i} \quad (9)$$

where

$D(x_1, x_2)$: Levensthein distance between two sequences
k_i : Number of actual journeys assigned to the representative journey i
S_i : Representative sequence i
S_{ij} : Sequence of actual journeys j assigned to i

[00218] An example metric for Internal Evaluation is *Distance gain [10]*. The distance gain measures the gain in using a representative journey instead of the true center of the set (i.e., the medoid of the whole dataset). In other words, it measures the gain obtained in using multiple representative journeys instead of a single one.

$$DistGain_i = \frac{\sum_{j=1}^{k_i} D(C(\sigma_{\mathcal{A}}), S_{ij}) - \sum_{j=1}^{k_i} D(S_i, S_{ij})}{\sum_{j=1}^{k_i} D(C(\sigma_{\mathcal{A}}), S_{ij})} \quad (10)$$

where

$D(x_1, x_2)$: Levensthein distance between two sequences
k_i : Number of actual journeys assigned to the representative journey i
S_i : Representative sequence i
S_{ij} : Sequence of actual journeys j assigned to i
$C(x)$: True center of the set

[00219] The platform 100 can evaluate two settings of the genetic algorithm compared to techniques that are used to cluster and summarize sets of sequential and categorical data.

[00220] Figure 22 depicts the approach at a high-level. As can be seen, with traditional approaches, a first step is to build a distance metric. This can use the edit distance with a constant cost operation set to 1. Once the distance matrix is built, this can involve creating k clusters. There can be a test using from 2 to 12 clusters and use the squared Calinski-Harabasz index to return the most statistically relevant. Then, we return the best representatives of each cluster using the neighborhood density, the centrality, the frequency, or the likelihood using the package Traminer. These techniques do not use the contextual data. Hence, to allow for a fair comparison, we compare these techniques with a version of our genetic algorithm that does not use contextual data. We call this version $Genetic_1$. We also test our genetic algorithm with a version that considers the contextual data, called $Genetic_2$. The detailed parameters we used are depicted in Table 3. Note that both the traditional and genetic approaches use the same techniques to find k_h (i.e., the squared Calinski-Harabasz index) and the distance is measured using the edit distance with a constant cost operation set to 1. Due to the non-deterministic nature of the genetic algorithm, we run it ten times for each setting.

[00221] The platform 100 can present the results using the external evaluation metrics, then the internal ones. Figure 23 shows that the solution that reduces the Jaccard distance the most is $Genetic_2$. This is not surprising, as this solution uses more information (i.e., the contextual data). Next, the best solution that does not use contextual data is $Genetic_1$. The drop in the average absolute distance in the number of journeys from the ground truth (which is visible in Figure 24) might be why the genetic algorithm outperforms the other techniques that solely rely on the Calinski-Harabasz index to determine the ideal number of journeys.

[00222] The internal evaluation of Figure 25 shows that not only does the genetic algorithm outperform the traditional approaches, it also proposes a better solution than the ground truth. This can be explained by the fact that when platform 100 injects noise, its potentially changes the solution.

[00223] Figure 26 summarizes the results.

EXPERIMENTS USING REAL DATASETS

[00224] This section reports on the experiments with a real dataset, the goal being to illustrate how a change in the settings impacts the results. Overall, the main settings are the same as

those described in the previous section. That is, the distance between sequences is measured using the edit distance with a constant cost operation set to 1, k_h is tested from 2 to 12 clusters using the squared Calinski-Harabasz index, p is 100, and e is 5. The parameters we adjust are the weight given to the fitness, the number of representatives, and the contextual distance.

[00225] We used a publicly available dataset describing the activities performed throughout the day by Chicago's citizens. There are 15 types of activities, such as, 'being at home', 'attending class', 'going shopping', or 'doing households errands'. In the context of this dataset, a journey is the sequence of activities starting from the morning until the night. Typically, 'being at home' → 'attending class' → 'being at home' is a journey consisting of three activities. The total number of journeys is 29,541 and there are 123,706 activities (with an average of 4.817 activities per journey). This dataset is interesting not only for the relatively large number of data points describing life trajectories, but also because of the available detailed contextual data, such as information on the citizens' demographics.

[00226] As a preprocessing step, the age of the customers can be binarized into five groups (see the legends in Figure 27) and other contextual data can be removed. This is an example. The goal is to show the influence of the settings on the journey map. Figure 27 shows the results using three different configurations.

[00227] In configuration 1, the contextual data is not leveraged (i.e., the contextual distance weight is set to 0). The platform 100 can interpret the resulting journey map as follows. The first journey represents people going to 'work', going back 'home' at noon, and returning to 'work' in the afternoon. The second journey is close to the first one, the main difference being that people do not seem to go back 'home' at noon. The third journey shows citizens being at 'home', going 'shopping' twice in the afternoon, and going back 'home'. Interestingly, the activity 'shopping' appears twice in a row. Note that this does not necessarily mean that the activity is repeated twice in a row; it might be that there are some other activities in between, but because they are not significant enough, they are not shown.

[00228] In configuration 2, platform 100 can test the effect on the resulting journey map when considering the ages of the customers. Therefore, we changed the weight put on the contextual distance from 0 to 1. As can be seen in Fig. 16, three representative journeys were generated. Each of these journeys has three touchpoints. They start from 'home' and finish at 'home'. In between, the first journey has the activity 'work', the second one has the activity 'shopping', and the last one the activity 'attending class'. It is interesting to note the effect of the configuration on

the contextual data (the distribution charts on the right side of Fig. 16). Indeed, while the age was equally distributed for each journey in configuration 1, we can observe that the age is discriminant in configuration 2. For instance, more than half of the citizens in the journey j_3 are under 16 years old, while this population represents only 8.7% of the entire dataset.

[00229] In configuration 3, platform 100 can show the effect when we increase the weight put on the contextual distance parameter. Journeys j_1 and j_3 are identical to those in configuration 2. However, a new and rather complex journey j_2 emerges. It starts from being at 'home', goes to 'health care', and then alternates between the activities 'entertainment' and 'home'. As can be expected, the effect of the contextual data is stronger than in configuration 2. We also observe that the distribution is impacted when giving more weight to homogeneity. We interpret the result as follows: Citizens younger than 29 years old tend to have two typical patterns of activities involving either 'school' or 'entertainment' while the most typical journeys for the other citizens involve 'work'.

[00230] Of course, this is a simplified overview of the data. For the almost 30,000 actual journeys in the event logs, there are numerous unique actual journeys that differ from the representative journeys we get from these three configurations. By letting the user choose the weight for each parameter, we let them explore different perspectives of the data. We claim that the best parameters depend on the dataset, the business context, and the goal of the exploration.

[00231] The platform 100 can use different quality criteria to guide the evolution process of discovering the best representative journeys for a given set of actual journeys that otherwise would be unreadable on a display device or interface. The platform 100 can perform well on synthetic datasets and that they allow us to discover journey maps showing different alternative perspectives of the same real dataset. However, quality criteria might be created to fit a business context. Typically, the criterion *simplicity* might be used in addition to or to replace k_h. Because it is difficult to anticipate what will be the ideal number of clusters returned by k_h, simplicity would offer a more intelligible way to choose between a highly fitting but more complex journey map, or a less fitting but simpler journey map. Industry-driven quality criteria might be used in addition to the ones proposed.

[00232] For the contextual distance, platform 100 can limit the set of contextual data to categorical data and–in the experimentation using a real dataset–platform 100 can select one item of contextual data manually (i.e., the age). The platform 100 can also preprocess the data to

distribute the age into bins (e.g., 60+ years old) to turn the continuous feature into a categorical one. Datasets can possess numerous items of contextual data of different types.

[00233] Finally, a genetic approach is one example solution to solve the customer journey discovery task.

[00234] The genetic approach to summarizing a set of customer journeys with the purpose of displaying them on a journey map offers an interesting alternative. First, the quality of the results is better, which is true using both internal and external evaluation metrics. Second, the weights of the three quality criteria are a flexible way to analyze a dataset under different perspectives. Third, in addition to the sequence of activities, our genetic algorithm can leverage contextual data to group similar journeys. By doing so, platform 100 can provide a way to summarize insights from customers that are hidden in the data.

[00235] The platform 100 can generate an interface to guide end-user in the process of making abstract traces from complex ones by leveraging the knowledge about the control-flow of activities stemmed from process models (produced by process mining discovery algorithm.

[00236] A trace or map is a conceptual tool used to visualize typical customers' trajectories when using a service. In their simplest form, maps show the interactions between a customer and a service provider through time. A series of interactions is called a journey. Because maps give a company a better understanding of their customers, they are becoming increasingly popular amongst practitioners.

[00237] A map can anticipate the best – or worst – journeys possible. Such journeys, displayed on a map, are called the expected journeys. However, customers might experience a different journey from the one anticipated. For this reason, platform 100 can leverage traces left by customers in information systems to build maps from evidence. Because the journeys that will be displayed on the map are produced from facts, they can be referred to as actual journeys.

[00238] However, when dealing with numerous journeys, it becomes unrealistic to display all the actual journeys on a single map at interface application 130. For illustration purposes, Figure 28 depicts 10,000 instances of the traces related to the handling of reviews for a journal, an example synthetic dataset. In the context of this dataset, the service provider is the conference's organizing committee, the customers are the researchers submitting their papers, and a journey describes the handling of the reviews, from the submission until the final decision. In Figure 28, at 2802 it is difficult to apprehend the typical paths of the reviewing process. To this end,

representative journeys have been introduced as a means of reducing the complexity. Indeed, the central map 2804 uses two representative journeys to summarize 10,000 actual journeys.

[00239] Although representative journeys decrease the complexity by reducing the number of journeys, a map might still be difficult to apprehend when it is composed of many activities. Indeed, even though only representative journeys are being used, quickly spotting the main differences between the two journeys visible in 2804 (Figure 28) is not straightforward due to the high number of activities and the length of the journeys.

[00240] The platform 100 can implement a map abstractor that leverages the expertise of process discovery to abstract maps for visual representation. More precisely, platform 100 can take as an input a process tree, parse it, starting from the leaves, and iteratively ask the end-user if it is relevant to merge the activities that belong to the same control-flow, and, if so, to provide a name for this group of activities. By doing so, platform 100 can let the end-user decide which activities should be merged and how they should be renamed. Then, one can visualize the same maps at different levels of granularity using a slider, which is visible in Figure 28 at 2806. At a certain level of granularity, platform 100 can enable an end user to observe, given the end activities, that one representative journey summarizes the accepted papers, while the other one depicts the rejected papers. The map abstractor can explore a seamless integration of business process models with journeys maps.

[00241] The platform 100 can implement process mining and process discovery activity. Section The platform 100 can implement customer journey discovery activity.

[00242] The platform 100 can implement integration of Process Mining with Journey Mapping and showcases the impact that the latter can have in the analysis of journeys. Process mining involving machine learning, data mining and process modeling and analysis. The platform 100 can implement the discovery of process models, one of the three types of process mining along with conformance and enhancement.

[00243] The idea behind the discovery of process models is to leverage the evidence left in information systems to build process models from event logs. The resulting process models are, therefore, based on factual data, showing how the process was really executed. To build such a model, process mining uses an input data format called event logs. An event log is a collection of traces, a trace being a single execution of a process composed of one or multiple activities.

[00244] For illustration purposes, let T = (⟨bdcef⟩; ⟨acdefg⟩; ⟨bcdefgg⟩) be an event log composed of 3 traces and 7 distinct activities. Regardless of the notation, the resulting models can express the control-flow relations between activities. For instance, for the event log, T, the model might express the following notation: 1) A and B are in an XOR relation (X); i.e., only one of them is executed; 2) C and D are executed in parallel (+); i.e., both activities are executed in any order; 3) E and F are in a sequence relation (→); i.e., F always follows E; 5) G is in a XOR loop (Combination of X and ↺); i.e., it can be executed 0 or many times. Note that T denotes a silent activity. It is used to correctly execute the process but it will not result in an activity which will be visible in the event logs. Figure 29 displays the five aforementioned relations using a process tree.

[00245] Discovering a process model from event logs is a challenge. The platform 100 can be robust enough to generalize (to avoid overfitting models) without being too generic. The platform 100 can build process models or process trees.

[00246] A process tree is an abstract hierarchical representation of a process model, where the leaves are annotated with activities and all the other nodes are annotated with operators such as X. Process trees are a means of guaranteeing the soundness of the discovered models. A model is considered to be not sound when some activities cannot be executed or when the end of the process cannot be reached.

[00247] The soundness guarantee is one reason that we choose the process tree notation. There are also three other reasons. First, process models in block structure achieve best performance in terms of fitness, precision, and complexity. Second, the hierarchical structure of process trees is ideal to derive multiple levels of granularity. Finally, process trees can be used by top-performing process model algorithms, such as the inductive miner or the Evolutionary Tree Miner [6].

[00248] The platform 100 can use a process mining based model that allows us to map a standard event log from process mining (i.e., XES) to store customer journeys, as an attempt to bring customer journeys and process mining closer together.

[00249] Discovering a set of representative journeys that best describe the actual journeys observed in the event logs is a challenge inspired by the process discovery challenge. However, instead of describing the control flows of activities using a business process model, the main trajectories (i.e., the representative journeys) can be shown using a journey map. It encompasses

example challenges: (1) choosing the number of representatives. Let *k* be this number of representative journeys used on a journey map; (2) grouping actual journeys in *k* clusters; and (3) for each *k*, finding a representative journey. The platform 100 can overcome them.

[00250] First, the number of representative journeys used to summarize the entire actual journeys needs to be set. Looking at 2802 from Figure 28, it is difficult to say how many representative journeys should be used to summarize the data. The platform 100 can use different ways to find the number of representative journeys. First, the number of representative journeys can be set manually. It can also be set using standard model selection techniques such as the Bayesian Information Criterion (BIC) penalty or the Calinski-Harabasz index, for example.

[00251] Once *k* has been defined, actual journeys should be split in *k* clusters and a representative journey per cluster must be found. One of the ways, is to first define a distance function between actual journeys, such as the edit distance, or shingles, and to build a distance matrix; then, to split the actual journeys in *k* groups using hierarchical clustering techniques. Next, the representative can be found using a frequent sequence mining algorithm, by counting the density of sequences in the neighborhood of each candidate sequence, by taking the most frequent sequences, or by taking the medoid. Instead of inferring the representative from the distance matrix, it is also possible to obtain it using statistical modeling. The platform 100 can employ an Expectation-Maximization algorithm on a mixture of *k* Markov models, and then for each Markov model the journey with the highest probability becomes the representative.

[00252] The platform 100 implements map abstractor using different steps to render a map at different levels of abstraction. The steps 300 are depicted in Figure 30. In the first step, the goal is to build a process tree given an event log. Next, using the same event log, the goal is to build a journey map.

[00253] The third step consists of parsing the tree obtained in step 1. To this aim, the platform 100 can use a script (stored in memory 108) which parses the process tree (i.e., XML file) and performs a reverse Breadth-first search; i.e., traversing the operators in the tree from the lowest ones to the root in a level-wise way. Let ℓ be the number of operators in the process tree. At each of the ℓ operators of the process tree, platform 100 offers the opportunity to the end-user to merge the leaves under the operator. If the user chooses to merge the activities, she should provide a new name and the operator virtually becomes a leaf. If the end-user chooses not to merge the activities, platform 100 can keep the leaves intact. If the answer is no, platform 100 can keep the activities separated at all levels of granularities, and we also disable the parents' steps. Indeed,

platform 100 can postulate that if a user does not want to merge two activities at a low level of granularity, it does not make sense to merge them later at a higher Level of granularity.

[00254] Finally, in step 4, platform 100 can transform the journey map at different levels of abstraction. Let λ be the number of abstractions which will be available for a journey map. It can be seen as the number of steps that will be included in the sliders visible in Figure 28 at 2806 Note that λ is equal to the number of times the end-user decides to merge the activities and that λ = ℓ when the end-user merges all the activities. Let *operator* λ be the λ^{th} operator to be merged. Let GetLevelAbstraction (*journey map, λ, pt*) be a function that returns a journey map at the λ^{th} level of abstraction. Algorithm 1 shows how the function Abstract is recursively called to get to the level of abstraction λ. The parameter removeSeqRepeats in Algorithm 1 in line 7 emphasizes that continuous sequence of activities that are to be replaced, will be replaced by only one instance of the new name given for this operator. For instance, if the journey is "AABCBAC", the leaves that are to be replaced, are "A" and "B" and the new name is "X", the journey will become "XCXC". This reduces the length of the journeys and, thus, increases the abstraction. The following shows operations for getting to the level of complexity.

```
Input  : cjm, customer journey map
         λ, level of abstraction
         pt, process tree annotated with merging decisions
Output : cjm_λ, cjm at the level of abstraction λ
1 Function GetLevelAbstraction(cjm, λ, pt)
2     for i ← 0 to λ do
3         cjm → Abstract(cjm, pt.operator_i)
4     return cjm
5 Function Abstract(cjm, op)
6     foreach journey in cjm do
7         journey.replace(op.leaves, op.new_name, removeSeqRepeats = True )
8     return cjm
```

[00255] This section provides a running example of the tool implemented by platform 100. The running example is based on synthetic event logs describing the handling of reviews for a journal referred to here. It contains 10,000 journeys and 236,360 activities.

[00256] In the first step, we obtained a process tree by using the inductive miner with default parameters. It results in the process tree 3100 visible in Figure 31.

[00257] In the second step, platform 100 can obtain a journey map by: (1) measuring the distance between actual journeys using the edit distance; (2) building a dendrogram using a hierarchical clustering algorithm; (3) finding *k* using the Calinski-Harabaz Score (k=2); (4) finding representative journeys which results in a journey map which is visible in 2804 (Fig. 28).

[00258] In the third step, platform 100 can parse the XML in javascript. To traverse the tree, platform 100 can be using a tree-like data structures. The order in which the operators are parsed is depicted in Fig. 31 (i.e., 'step'). Fig. 31 shows that we decided to merge 7 out of the 9 operators. Note that we decided not to merge the activities 'reject' and 'accept', which disabled the option of merging all the activities below step 9. Figure 32 shows a screen capture 3200 of the application when merging the activities during step 1.

[00259] Finally, Figure 33 shows the resulting journey maps at three levels of abstraction 3302, 3304, 3306.

[00260] Journey maps are being used more and more to help service providers put themselves in their customers' shoes. The platform 100 provides automated ways of building them and displaying complex paths on a journey map. By leveraging process trees – a format build within the process mining community – platform 100 can bring customer journey analytics and process mining closer together.

[00261] The following discussion provides many example embodiments of the inventive subject matter. Although each embodiment represents a single combination of inventive elements, the inventive subject matter is considered to include all possible combinations of the disclosed elements. Thus if one embodiment comprises elements A, B, and C, and a second embodiment comprises elements B and D, then the inventive subject matter is also considered to include other remaining combinations of A, B, C, or D, even if not explicitly disclosed.

[00262] The embodiments of the devices, systems and methods described herein may be implemented in a combination of both hardware and software. These embodiments may be implemented on programmable computers, each computer including at least one processor, a data storage system (including volatile memory or non-volatile memory or other data storage elements or a combination thereof), and at least one communication interface.

[00263] Program code is applied to input data to perform the functions described herein and to generate output information. The output information is applied to one or more output devices. In some embodiments, the communication interface may be a network communication interface. In

embodiments in which elements may be combined, the communication interface may be a software communication interface, such as those for inter-process communication. In still other embodiments, there may be a combination of communication interfaces implemented as hardware, software, and combination thereof.

[00264] Throughout the foregoing discussion, numerous references will be made regarding servers, services, interfaces, portals, platforms, or other systems formed from computing devices. It should be appreciated that the use of such terms is deemed to represent one or more computing devices having at least one processor configured to execute software instructions stored on a computer readable tangible, non-transitory medium. For example, a server can include one or more computers operating as a web server, database server, or other type of computer server in a manner to fulfill described roles, responsibilities, or functions.

[00265] The technical solution of embodiments may be in the form of a software product. The software product may be stored in a non-volatile or non-transitory storage medium, which can be a compact disk read-only memory (CD-ROM), a USB flash disk, or a removable hard disk. The software product includes a number of instructions that enable a computer device (personal computer, server, or network device) to execute the methods provided by the embodiments.

[00266] The embodiments described herein are implemented by physical computer hardware, including computing devices, servers, receivers, transmitters, processors, memory, displays, and networks. The embodiments described herein provide useful physical machines and particularly configured computer hardware arrangements.

[00267] Although the embodiments have been described in detail, it should be understood that various changes, substitutions and alterations can be made herein.

[00268] Moreover, the scope of the present application is not intended to be limited to the particular embodiments of the process, machine, manufacture, composition of matter, means, methods and steps described in the specification.

[00269] As can be understood, the examples described above and illustrated are intended to be exemplary only.

WHAT IS CLAIMED IS:

Any and all features of novelty or inventive step described, suggested, referred to, exemplified, or shown herein, including but not limited to processes, systems, devices, and computer-readable and –executable programming and/or other instruction sets suitable for use in implementing such features.

1. A platform comprising:

a data storage device for storing event traces, the event traces having attributes that indicate activities over time;

a processor configured to process machine executable instructions to generate visual elements for an interactive interface application by:

computing clusters of the event traces based on a genetic process function, each cluster corresponding to a set of event traces and a representative event trace based on representative attributes of the set of event traces, the genetic process function mapping each event trace to a cluster;

generating the visual elements for an interface application, the visual elements indicating the clusters and, for each cluster, the representative event trace based on the representative attributes, the representative event trace summarizing the set of event traces corresponding to the cluster;

controlling rendering of the interface application at a device to display the visual elements and a plurality of selectable portions; and

responsive to a selection of a selectable portion of the plurality of selectable portions, controlling rendering of the interface application at the device to navigate the clusters.

2. The platform of claim 1 wherein the processor is configured to implement the genetic process function by: evaluating initial set representative event traces to generate elite set representative event traces, generating additional initial representative event traces using a transformation process, evaluating the additional initial representative event traces to generate additional elite set representative event traces, continuing the generating and

the evaluating until stopping criterion is met, using resulting elite representative event traces as the representative event traces for the clusters.

3. The platform of claim 1 wherein, for each event trace, the attributes are represented as a two-dimensional tuple of activities over time and contextual data and the genetic process function can evaluate the two-dimensional tuples of the event traces.

4. The platform of claim 1 wherein the processor computes, for each cluster, the representative event trace as a pattern of activities to summarize the patterns contained within the set of event traces of the cluster.

5. The platform of claim 1 wherein the processor implements the genetic process function by pre-processing the event traces, determining an initial population that corresponds to a number of the clusters, evaluating initial representative event traces, evaluating stop criteria, implementing a genetic operation to generate additional initial representative event traces, outputting a map abstraction for use in generating additional visual elements.

6. The platform of claim 5 wherein the processor is configured to evaluate the initial representative event traces based on quality criteria for fitness, a number of representatives, a contextual distance and an average quality.

7. The platform of claim 6 wherein the processor is configured to weigh the quality criteria to generate an overall quality and average the overall quality.

8. The platform of claim 1 wherein the processor is configured to evaluate the initial representative event traces based on internal and external evaluation metrics.

9. The platform of claim 1 wherein the processor is configured to:

generate visual elements for multiple views for the interface application, the multiple views having a first view indicating a pattern of activities for the set of representative attributes; a second view indicating the cluster and the pattern of activities within the cluster; and a third view indicating descriptors for the cluster, the multiple views having a plurality of selectable portions;

controlling rendering of the interface application at a device to display the multiple views synchronously; and

responsive to a selection of a selectable portion of the plurality of selectable portions, controlling rendering of the interface application at the device to update the multiple views synchronously based on the selected portion to navigate the cluster.

10. A non-transitory computer readable medium storing machine executable instructions to configure a processor to:

compute clusters of the event traces based on a genetic process function, each cluster corresponding to a set of event traces and a representative event trace based on representative attributes of the set of event traces, the genetic process function mapping each event trace to a cluster;

generate the visual elements for an interface application, the visual elements indicating the clusters and, for each cluster, the representative event trace based on the representative attributes, the representative event trace summarizing the set of event traces corresponding to the cluster;

control rendering of the interface application at a device to display the visual elements and a plurality of selectable portions; and

responsive to a selection of a selectable portion of the plurality of selectable portions, control rendering of the interface application at the device to navigate the clusters.

11. The non-transitory computer readable medium of claim 10 wherein the machine executable instructions configure the processor to implement the genetic process function by: evaluating initial set representative event traces to generate elite set representative event traces, generating additional initial representative event traces using a transformation process, evaluating the additional initial representative event traces to generate additional elite set representative event traces, continuing the generating and the evaluating until stopping criterion is met, using resulting elite representative event traces as the representative event traces for the clusters.

12. The non-transitory computer readable medium of claim 10 wherein, for each event trace, the attributes are represented as a two-dimensional tuple of activities over time and contextual data and the genetic process function can evaluate the two-dimensional tuples of the event traces.

13. The non-transitory computer readable medium of claim 10 wherein the machine executable instructions configure the processor to compute, for each cluster, the representative event trace as a pattern of activities to summarize the patterns contained within the set of event traces of the cluster.

14. The non-transitory computer readable medium of claim 10 wherein the machine executable instructions configure the processor to implement the genetic process function by pre-processing the event traces, determining an initial population that corresponds to a number of the clusters, evaluating initial representative event traces, evaluating stop criteria, implementing a genetic operation to generate additional initial representative event traces, outputting a map abstraction for use in generating additional visual elements.

15. The non-transitory computer readable medium of claim 10 wherein the machine executable instructions configure the processor to evaluate the initial representative event traces based on quality criteria for fitness, a number of representatives, a contextual distance and an average quality.

16. The non-transitory computer readable medium of claim 10 wherein the machine executable instructions configure the processor to weigh the quality criteria to generate an overall quality and average the overall quality.

17. The non-transitory computer readable medium of claim 10 wherein the machine executable instructions configure the processor to evaluate the initial representative event traces based on internal and external evaluation metrics.

18. The non-transitory computer readable medium of claim 10 wherein the machine executable instructions configure the processor to:

generate visual elements for multiple views for the interface application, the multiple views having a first view indicating a pattern of activities for the set of representative attributes; a second view indicating the cluster and the pattern of

activities within the cluster; and a third view indicating descriptors for the cluster, the multiple views having a plurality of selectable portions;

controlling rendering of the interface application at a device to display the multiple views synchronously; and

responsive to a selection of a selectable portion of the plurality of selectable portions, controlling rendering of the interface application at the device to update the multiple views synchronously based on the selected portion to navigate the cluster.

19. A platform comprising:

a data storage device for storing event traces, the event traces having attributes that indicate activities over time, the event traces relating to current data and historic data;

a processor configured to process machine executable instructions to generate visual elements for an interactive interface application by:

generating a mapping of process mining elements to the event traces, the mapping having an XML-based set of concepts, the mapping linking process de facto models and de jure models to actual and expected customer experiences;

generating a data structure for hierarchical components of a journey map, each component having an element and an attribute, the journey map having at least one element being a root node of the hierarchical components, the actual and expected customer experiences corresponding to the journey map;

generating a visual elements for the components of the journey map; and

transmitting the visual elements to an interface application for display on a device.

20. A platform comprising:

a data storage device for storing event traces, the event traces having attributes that indicate activities over time;

a processor configured to process machine executable instructions to generate visual elements for a map abstractor of an interactive interface application by:

generating a process tree for the event traces by grouping the event traces based on similar attributes, the process tree having nodes corresponding to the event traces, the process tree having leaf nodes;

parsing the process tree;

starting from leaf nodes of the parsed process tree, iteratively generating a prompt at an interface application to merge a set of event traces that belong to a subset of the parsed process tree;

upon receiving confirmation to merge the set of event traces, generating a name for the set of event traces and replacing the nodes corresponding to the set of event traces with a new node indicating the name and a set of representative attributes for the set of event traces;

generating the visual elements for an interface application, the visual elements indicating the process tree and the new nodes, the visual elements representing an abstracted process tree; and

controlling rendering of the interface application at a device to display the visual elements.

21. The platform of claim 10 wherein a process tree is an abstract hierarchical representation of a process model, where the leaves are annotated with activities and all the other nodes are annotated with operators.

22. A platform comprising:

a data storage device having a data warehouse model for storing event traces, each event trace having attributes that indicate activities over time;

a processor configured to process machine executable instructions to generate visual elements for an interactive interface application by:

generating hierarchical cluster for the event traces by grouping event traces having similar attributes;

generating a set of representative attributes for a set of event traces of the hierarchical cluster, the set of representative attributes computed based on the similar attributes;

generating the visual elements for multiple views for the interface application, the multiple views having a first view indicating a pattern of activities for the set of representative attributes; a second view indicating the hierarchical cluster and the pattern of activities within the hierarchical cluster; and a third view indicating descriptors for the hierarchical cluster, the multiple views having a plurality of selectable portions;

controlling rendering of the interface application at a device to display the multiple views synchronously; and

responsive to a selection of a selectable portion of the plurality of selectable portions, controlling rendering of the interface application at the device to update the multiple views synchronously based on the selected portion to navigate the hierarchical cluster.

23. The platform of claim 22 wherein the processor generates the hierarchical cluster by computing a distance measure for the attributes of the event traces to compute the similar attributes.

24. The platform of claim 22 wherein the processor generates the hierarchical cluster by computing a distance measure based on a Levenshtein distance.

25. The platform of claim 22 wherein the processor generates the hierarchical cluster by computing a distance between clusters of the hierarchical cluster.

26. The platform of claim 22 wherein the pattern of activities indicates the set of representative attributes for a segment of the hierarchical cluster that corresponds to the set of event traces.

27. The platform of claim 26 wherein the second view indicates the segment of the hierarchical cluster for the pattern of activities indicated in the first view.

28. The platform of claim 22 wherein the selection portion is within a view of the multiple views to trigger an update to the other views of the multiple views.

29. The platform of claim 22 wherein the processor updates the multiple views synchronously based on the selected portion to navigate the hierarchical cluster at different levels of abstraction or granularity.

30. The platform of claim 22 wherein the processor generates the hierarchical cluster based on a proximity measure using an order of activities for the event traces.

31. The platform of claim 22 wherein the processor computes salient characteristics for the descriptors to indicate indexes based on a chi-square test.

32. The platform of claim 22 wherein the processor receives a navigation goal at the interface application, the goal indicating an attribute and responsive to the navigation goal, controlling rendering of the interface application at the device to update the multiple views synchronously to indicate data based on the navigation goal.

33. The platform of claim 22 wherein the processor generates the hierarchical cluster based on layers, a layer corresponding to a number of event traces that will be grouped based on the similar attributes, the layers corresponding to a height of the hierarchical cluster.

34. The platform of claim 22 wherein the processor generates the set of representative attributes based on the similar attributes and a frequent sequences mining process.

35. A non-transitory computer readable medium storing machine executable instructions to configure a processor to:

generate hierarchical cluster for the event traces by grouping event traces having similar attributes, each event trace having attributes that indicate activities over time;

generate a set of representative attributes for a set of event traces of the hierarchical cluster, the set of representative attributes computed based on the similar attributes;

generate the visual elements for multiple views for the interface application, the multiple views having a first view indicating a pattern of activities for the set of representative attributes; a second view indicating the hierarchical cluster and the pattern of activities within the hierarchical cluster; and a third view indicating descriptors for the hierarchical cluster, the multiple views having a plurality of selectable portions;

control rendering of the interface application at a device to display the multiple views synchronously; and

responsive to a selection of a selectable portion of the plurality of selectable portions, control rendering of the interface application at the device to update the multiple views synchronously based on the selected portion to navigate the hierarchical cluster.

36. The non-transitory computer readable medium of claim 35 wherein the machine executable instructions configure the processor to generate the hierarchical cluster by computing a distance measure for the attributes of the event traces to compute the similar attributes.

37. The non-transitory computer readable medium of claim 35 wherein the machine executable instructions configure the processor to generate the hierarchical cluster by computing a distance between clusters of the hierarchical cluster.

38. The non-transitory computer readable medium of claim 35 wherein the machine executable instructions configure the processor to update the multiple views synchronously based on the selected portion to navigate the hierarchical cluster at different levels of abstraction or granularity.

39. The non-transitory computer readable medium of claim 35 wherein the machine executable instructions configure the processor to receive a navigation goal at the interface application, the goal indicating an attribute and responsive to the navigation goal, controlling rendering of the interface application at the device to update the multiple views synchronously to indicate data based on the navigation goal.

40. The non-transitory computer readable medium of claim 35 wherein the machine executable instructions configure the processor to generate the hierarchical cluster based on layers, a layer corresponding to a number of event traces that will be grouped based on the similar attributes, the layers corresponding to a height of the hierarchical cluster.

41. A computer process to generate visual elements for an interactive interface application, the process comprising:

at a processor,

generating hierarchical cluster for the event traces by grouping event traces based on distance measures, each event trace having attributes that indicate activities over time;

generating a set of representative attributes for a set of event traces of the hierarchical cluster, the set of representative attributes computed based on the grouped event traces;

generating the visual elements for multiple views for the interface application, the multiple views having a first view indicating a pattern of activities for the set of representative attributes; a second view indicating the hierarchical cluster and the pattern of activities within the hierarchical cluster; and a third view indicating descriptors for the hierarchical cluster, the multiple views having a plurality of selectable portions;

controlling rendering of the interface application at a device to display the multiple views synchronously; and

responsive to a selection of a selectable portion of the plurality of selectable portions, controlling rendering of the interface application at the device to update the multiple views synchronously based on the selected portion, wherein the selection portion is within a view of the multiple views to trigger an update to the other views of the multiple views.

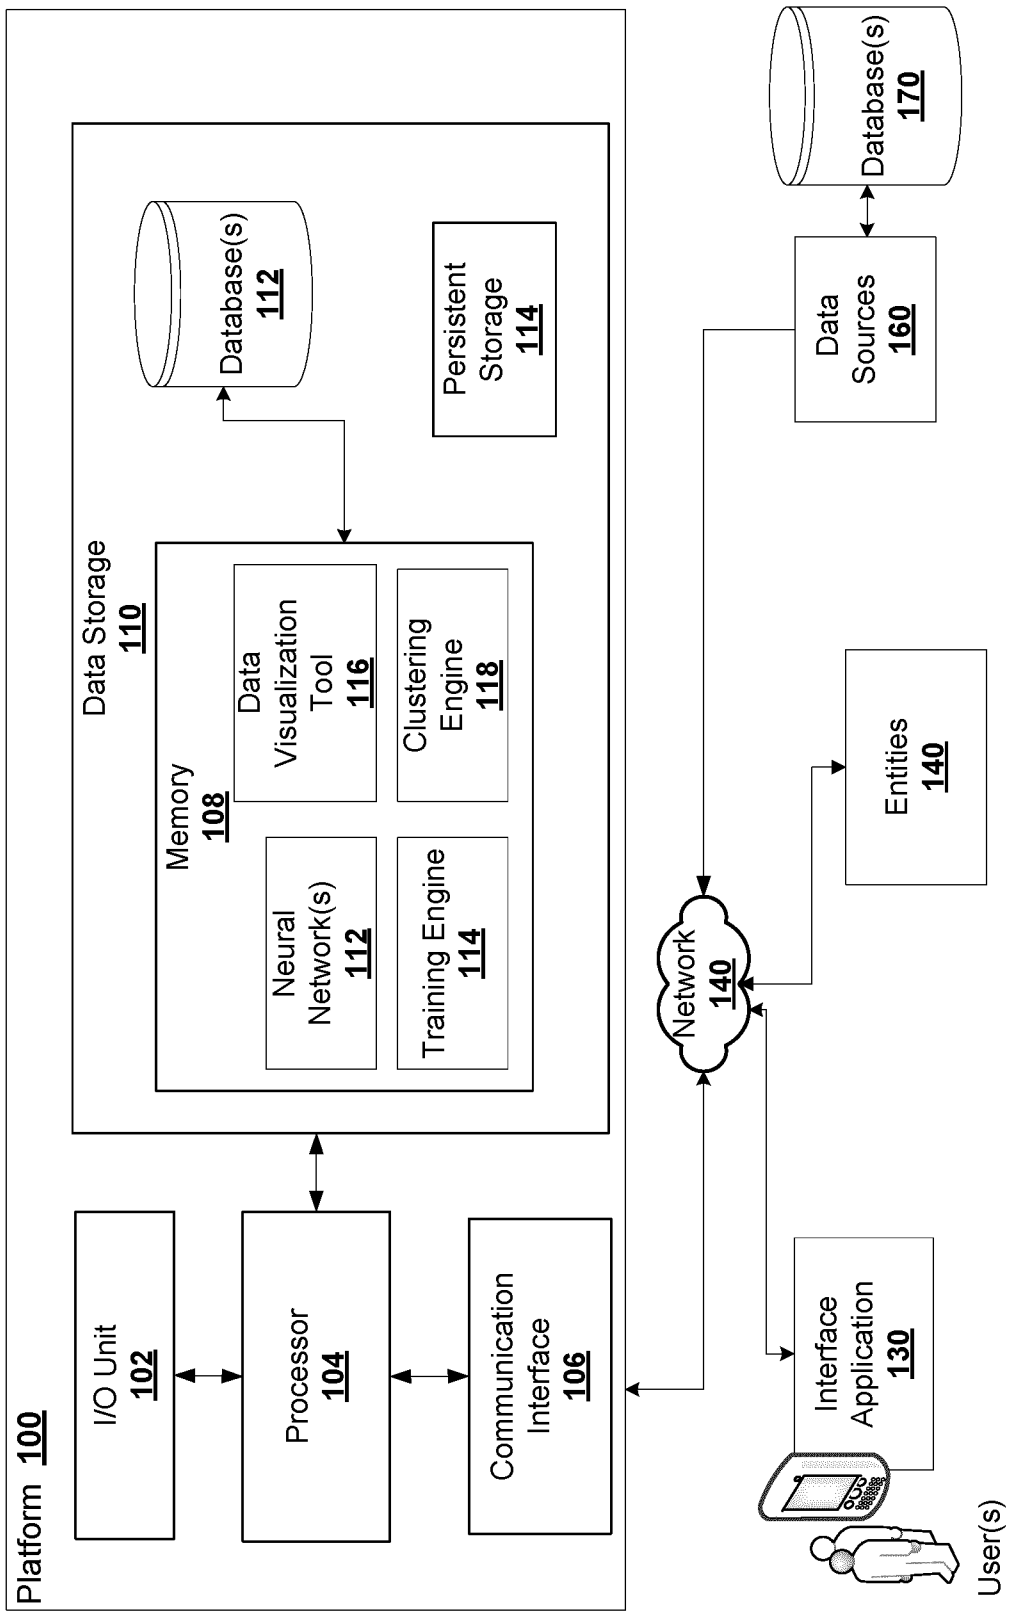
Platform 100
I/O Unit 102
Processor 104
Communication Interface 106
Data Storage 110
Memory 108
Data Visualization Tool 116
Clustering Engine 118
Neural Network(s) 112
Training Engine 114
Database(s) 112
Persistent Storage 114
Network 140
Data Sources 160
Database(s) 170
Entities 140
Interface Application 130
User(s)

FIG. 1

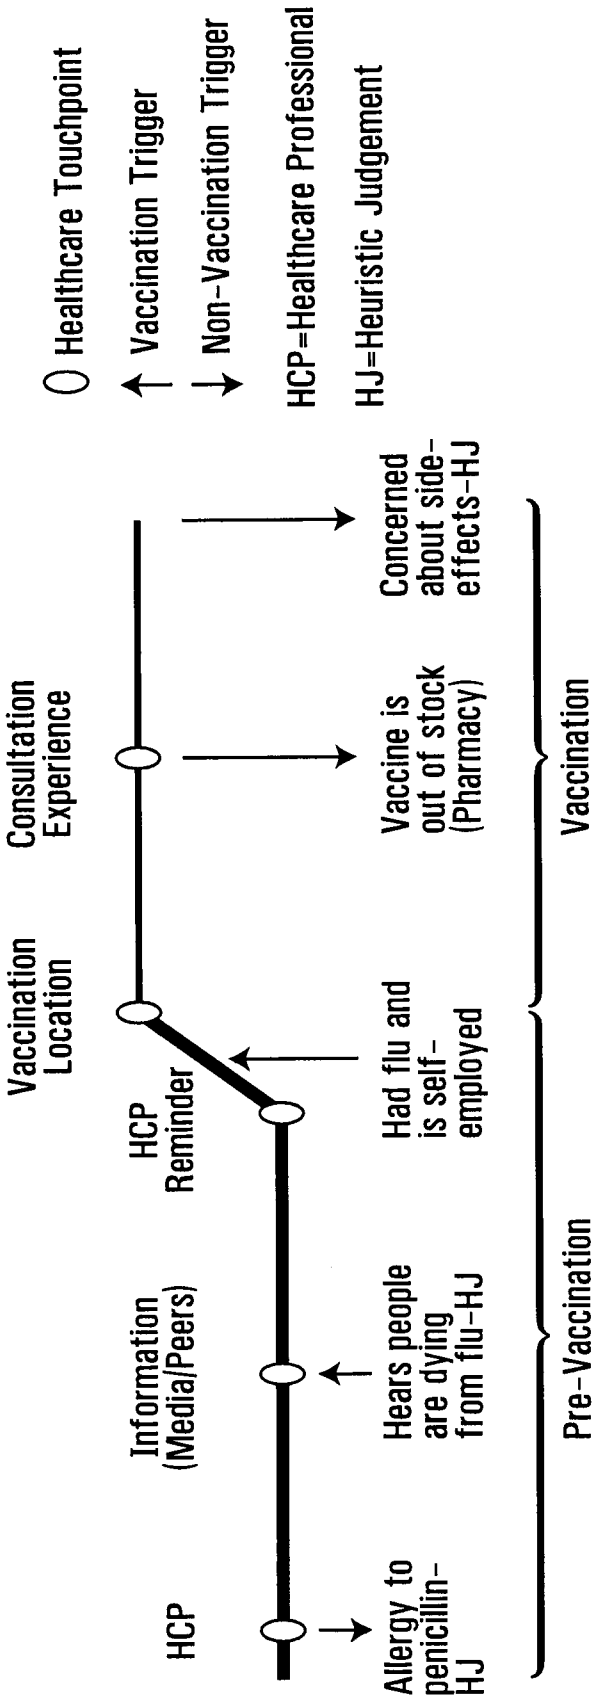
Healthcare Touchpoint
Vaccination Trigger
Non-Vaccination Trigger
HCP=Healthcare Professional
HJ=Heuristic Judgement
Vaccination Location
Consultation Experience
HCP
Information (Media/Peers)
HCP Reminder
Allergy to penicillin-HJ
Hears people are dying from flu-HJ
Had flu and is self-employed
Vaccine is out of stock (Pharmacy)
Concerned about side-effects-HJ
Pre-Vaccination
Vaccination

FIG. 2

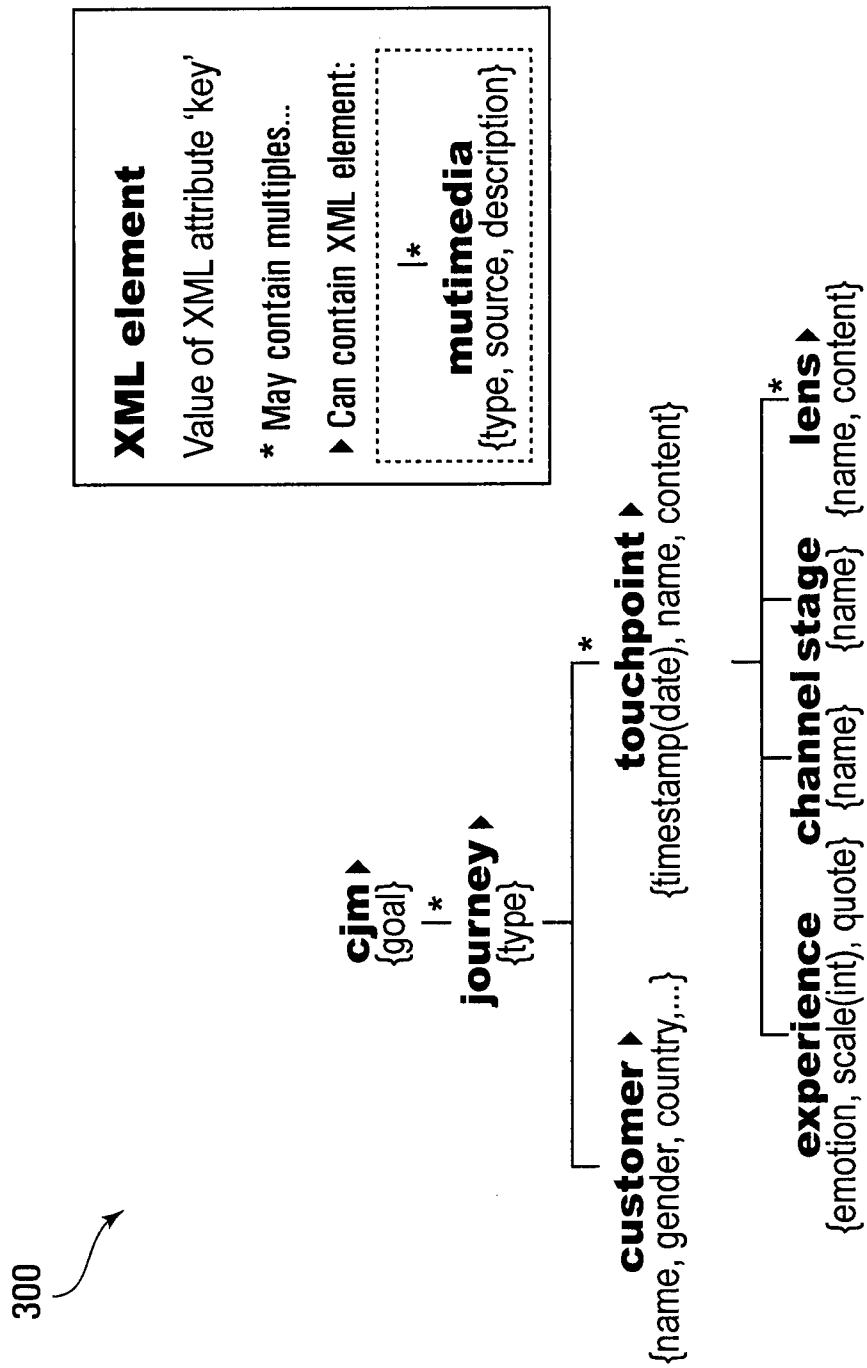
XML element
Value of XML attribute 'key'
* May contain multiples...
▸ Can contain XML element:
|*
mutimedia
{type, source, description}
300
cjm ▸
{goal}
|*
journey ▸
{type}
customer ▸
{name, gender, country,...}
*
touchpoint ▸
{timestamp(date), name, content}
experience
{emotion, scale(int), quote}
channel stage
{name}
{name}
*
lens ▸
{name, content}

FIG. 3

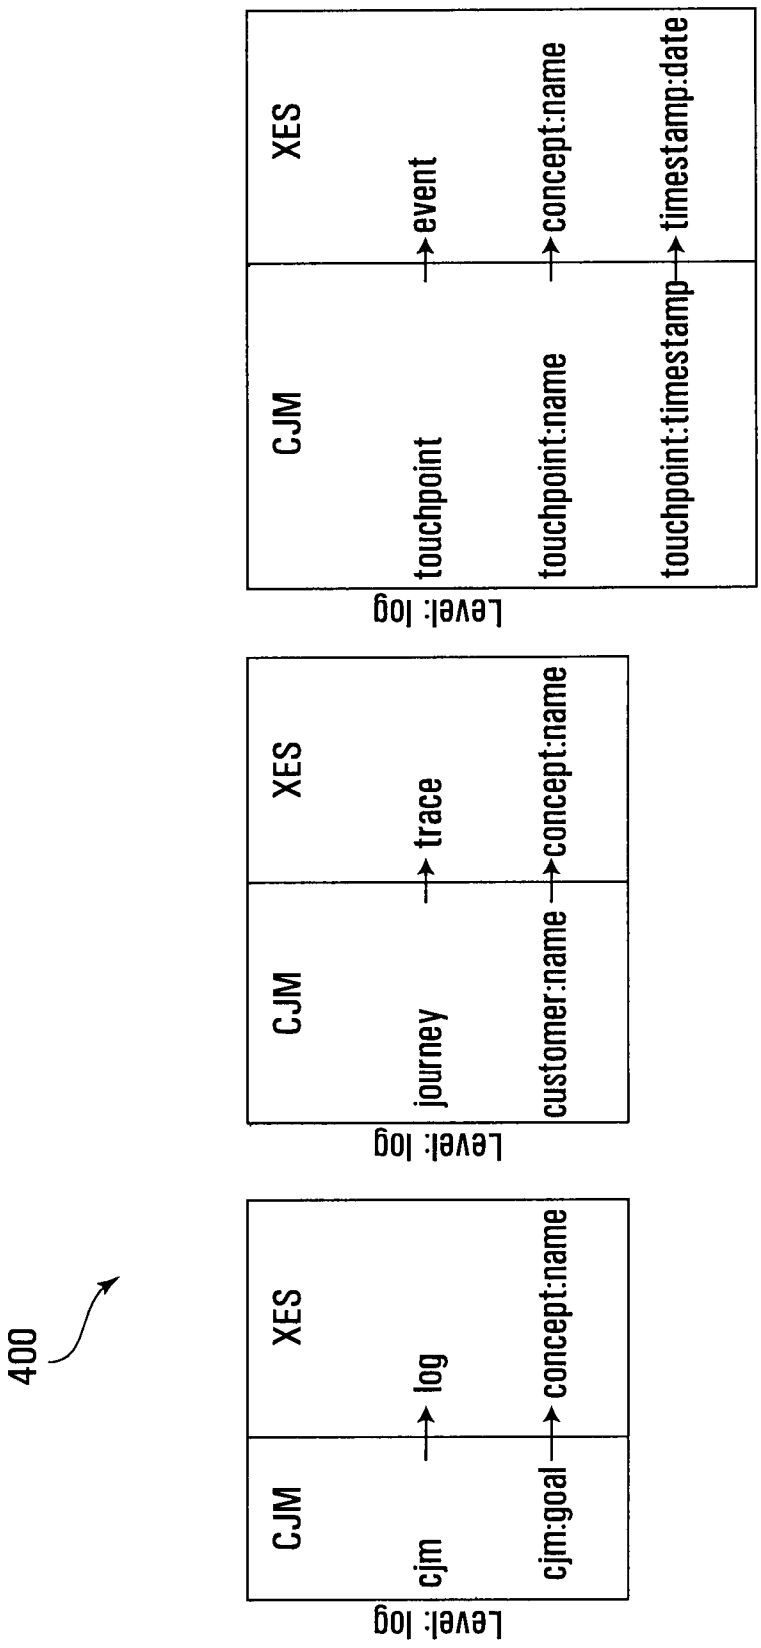
400
Level: log
CJM
cjm
cjm:goal
XES
log
concept:name
Level: log
CJM
journey
customer:name
XES
trace
concept:name
Level: log
CJM
touchpoint
touchpoint:name
touchpoint:timestamp
XES
event
concept:name
timestamp:date

FIG. 4

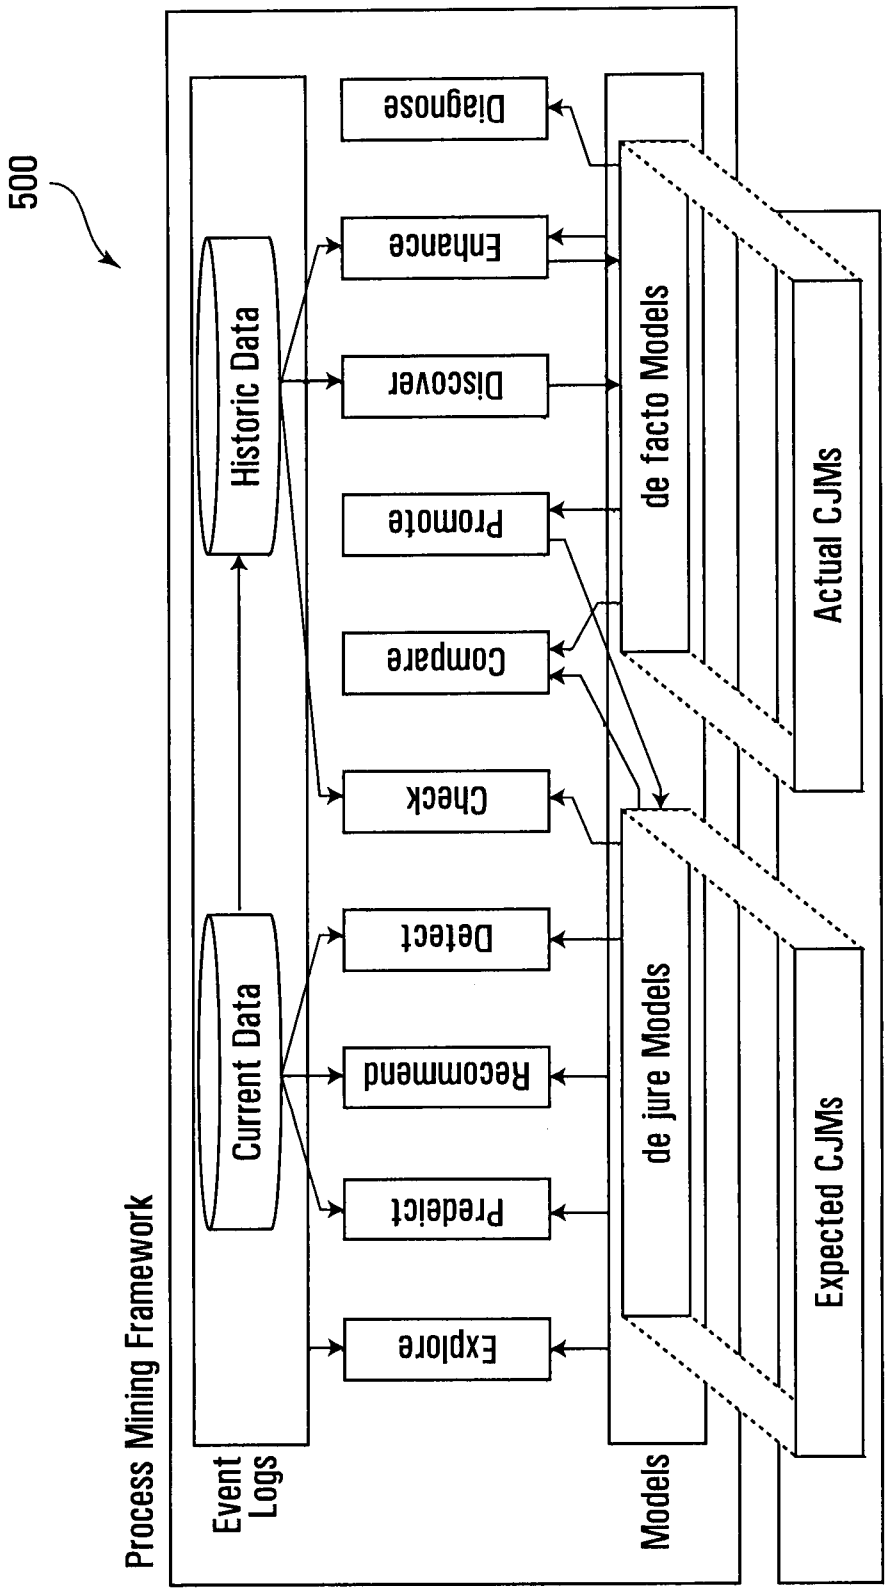
500
Process Mining Framework
Event Logs
Historic Data
Current Data
Diagnose
Enhance
Discover
Promote
Compare
Check
Detect
Recommend
Predict
Explore
Models
de facto Models
de jure Models
Customer Journey Mapping Extension
Actual CJMs
Expected CJMs

FIG. 5

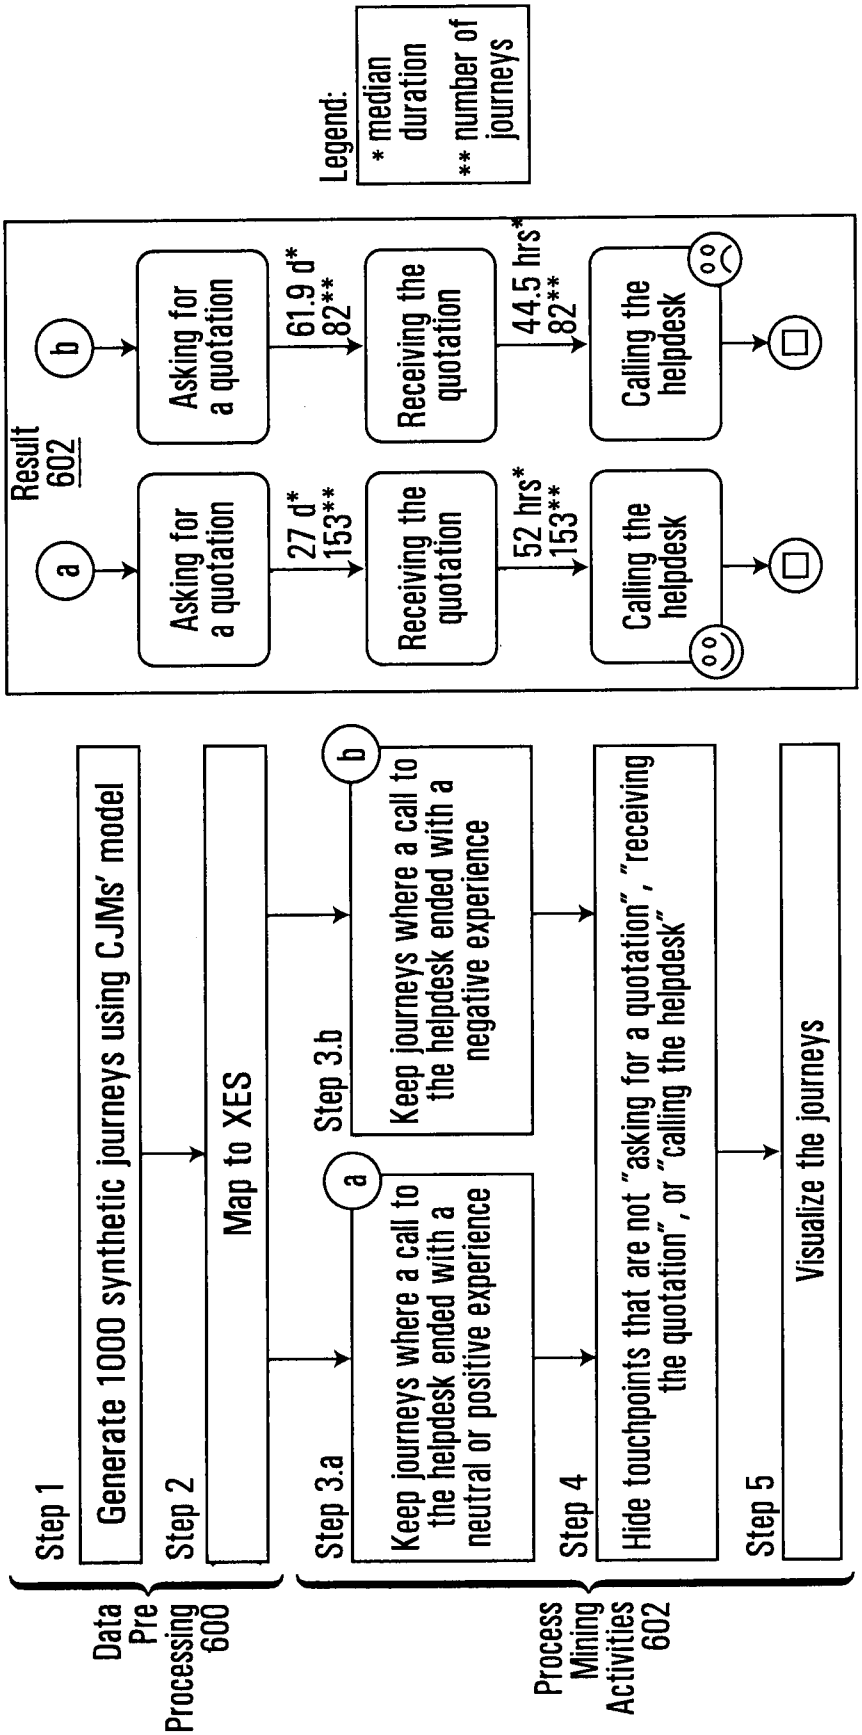
Legend:
* median duration
** number of journeys
Result 602
a
Asking for a quotation
27 d*
153**
Receiving the quotation
52 hrs*
153**
Calling the helpdesk
b
Asking for a quotation
61.9 d*
82**
Receiving the quotation
44.5 hrs*
82**
Calling the helpdesk
Data Pre Processing 600
Step 1
Generate 1000 synthetic journeys using CJMs' model
Step 2
Map to XES
Process Mining Activities 602
Step 3.a
a
Keep journeys where a call to the helpdesk ended with a neutral or positive experience
Step 3.b
b
Keep journeys where a call to the helpdesk ended with a negative experience
Step 4
Hide touchpoints that are not "asking for a quotation", "receiving the quotation", or "calling the helpdesk"
Step 5
Visualize the journeys

FIG. 6

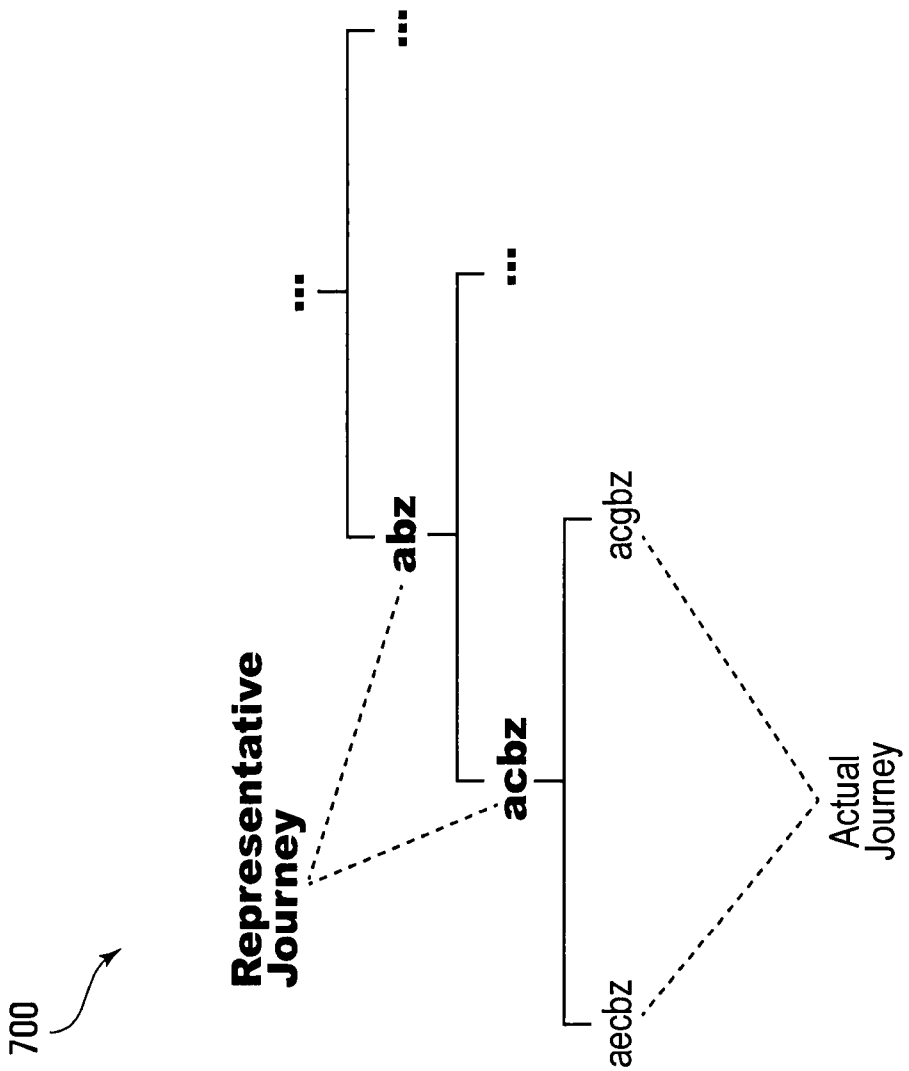
700
Representative Journey
abz
acbz
...
...
...
acgbz
aecbz
Actual Journey

FIG. 7

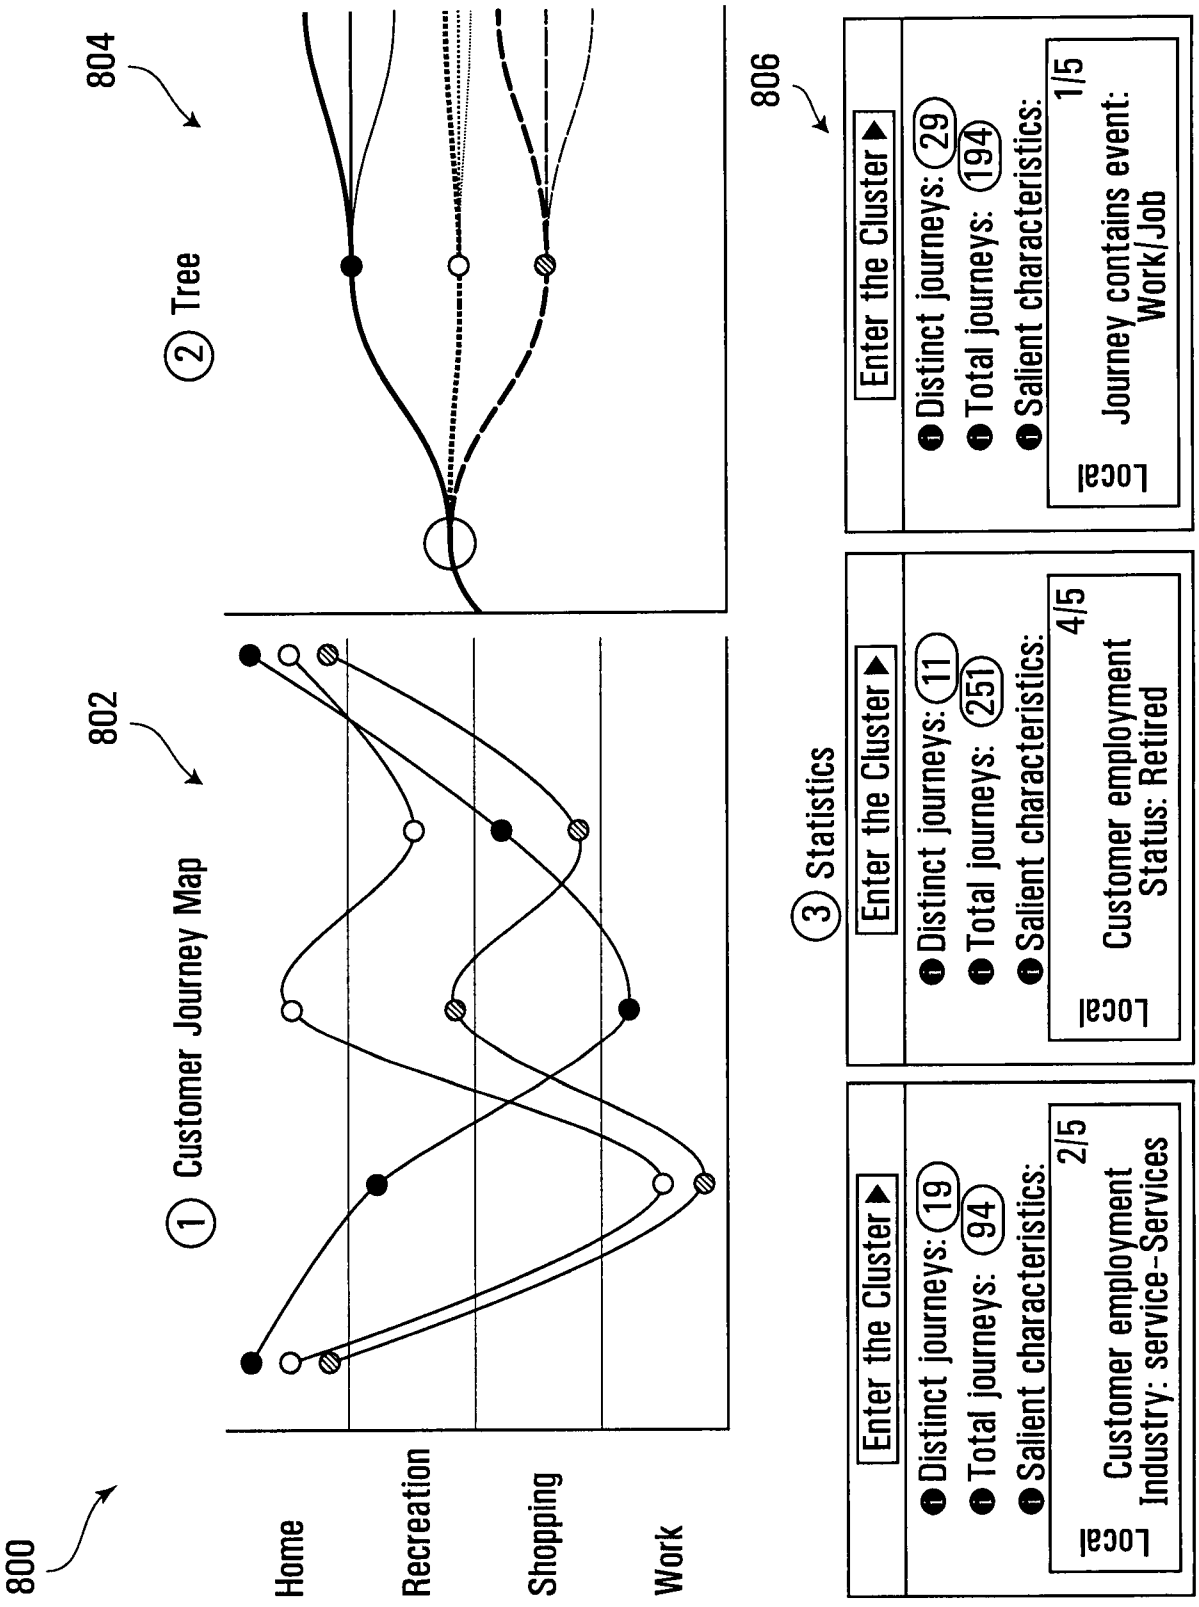
800
802
804
806
1 Customer Journey Map
2 Tree
3 Statistics
Home
Recreation
Shopping
Work
Enter the Cluster
Distinct journeys: 19
Total journeys: 94
Salient characteristics: 2/5
Customer employment Industry: service-Services
Local
Enter the Cluster
Distinct journeys: 11
Total journeys: 251
Salient characteristics: 4/5
Customer employment Status: Retired
Local
Enter the Cluster
Distinct journeys: 29
Total journeys: 194
Salient characteristics: 1/5
Journey contains event: Work/Job
Local

FIG. 8

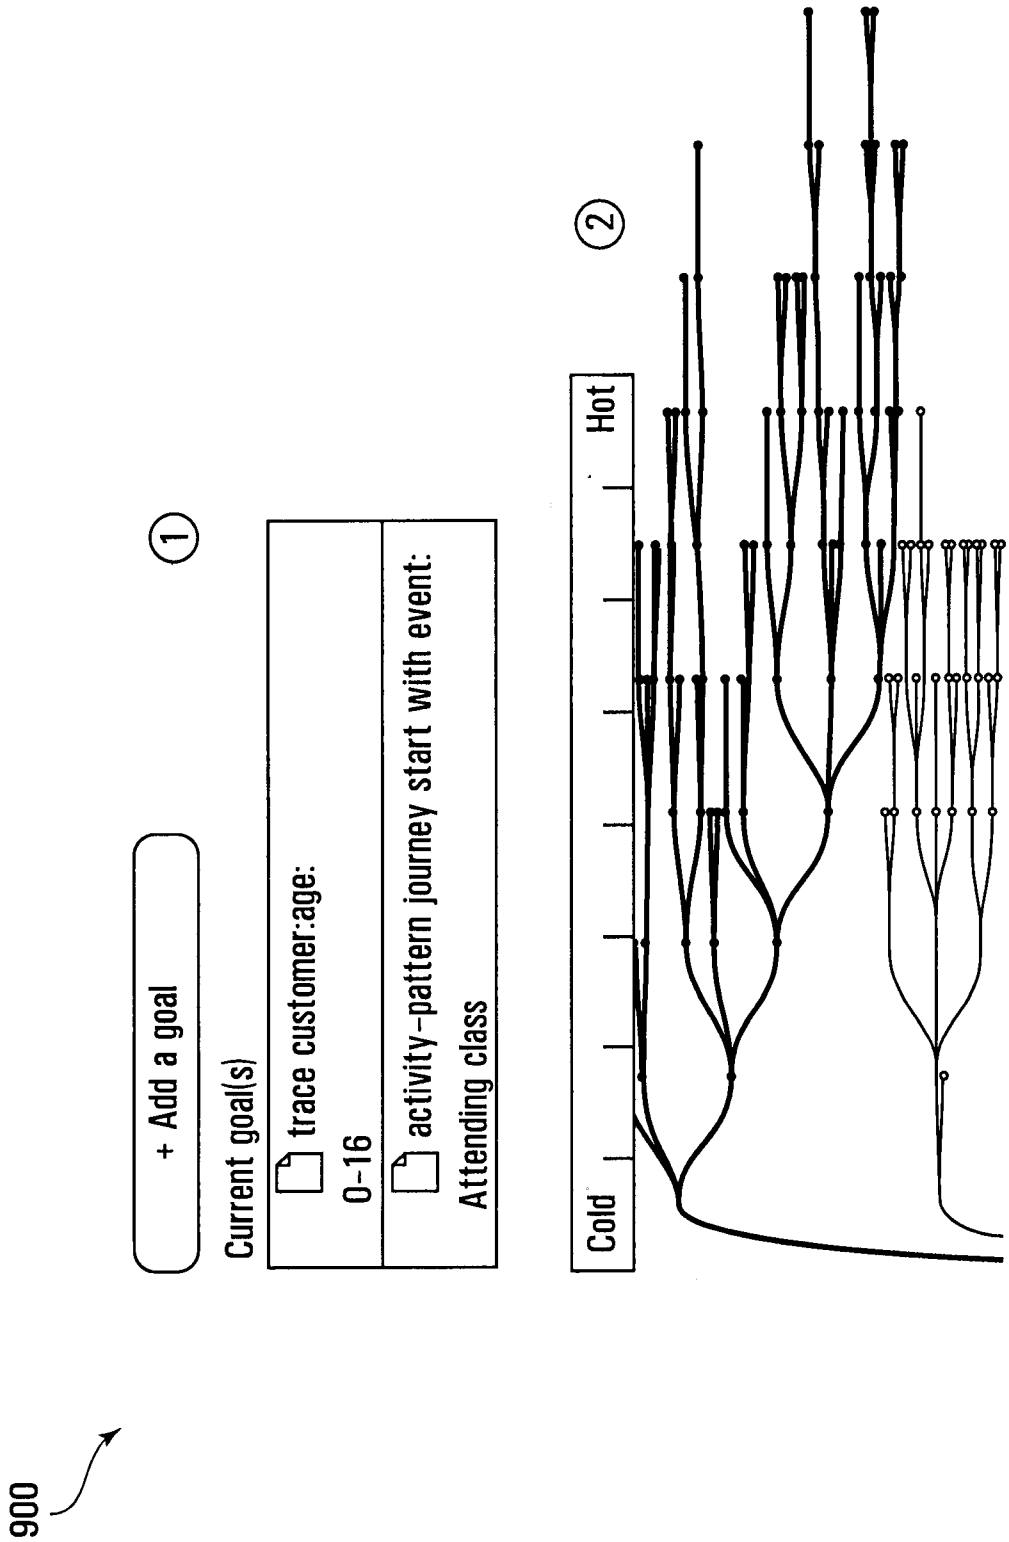
900
+ Add a goal
1
Current goal(s)
trace customer:age:
0-16
activity-pattern journey start with event:
Attending class
2
Hot
Cold

FIG. 9

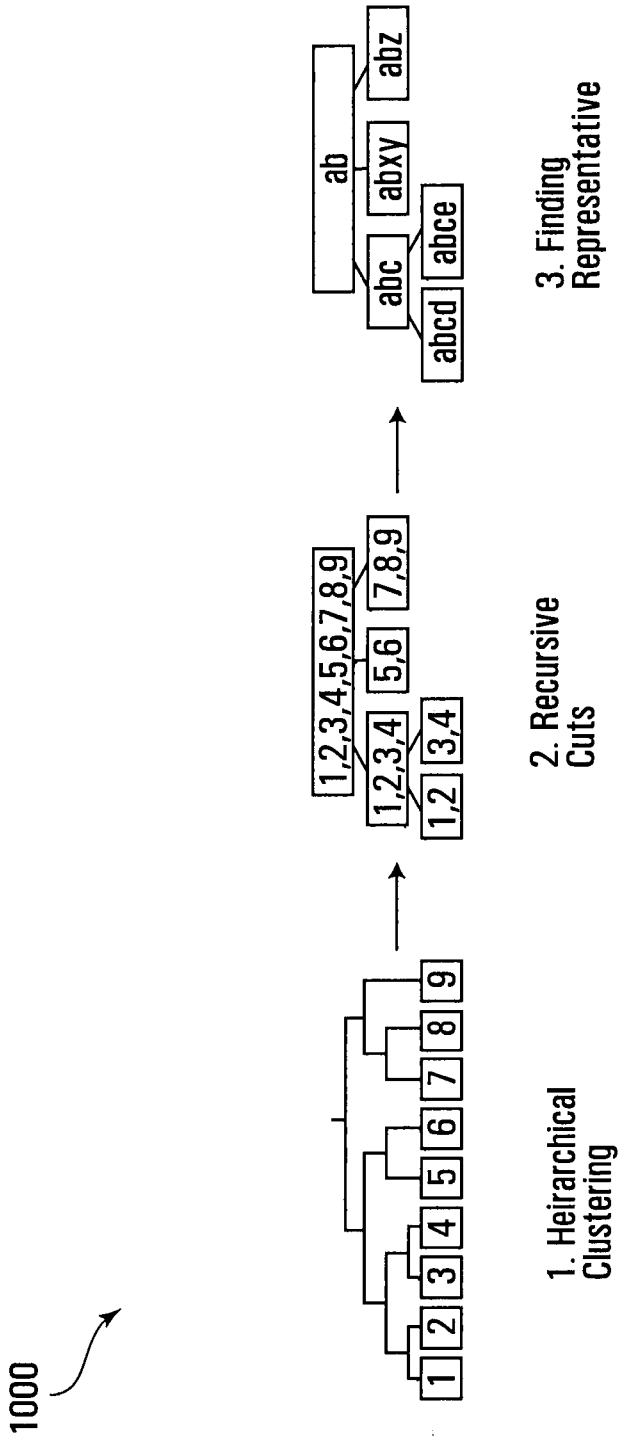
1000
1
2
3
4
5
6
7
8
9
1. Heirarchical Clustering
1,2
3,4
1,2,3,4
5,6
7,8,9
1,2,3,4,5,6,7,8,9
2. Recursive Cuts
abc
abcd
abce
abxy
abz
ab
3. Finding Representative

FIG. 10

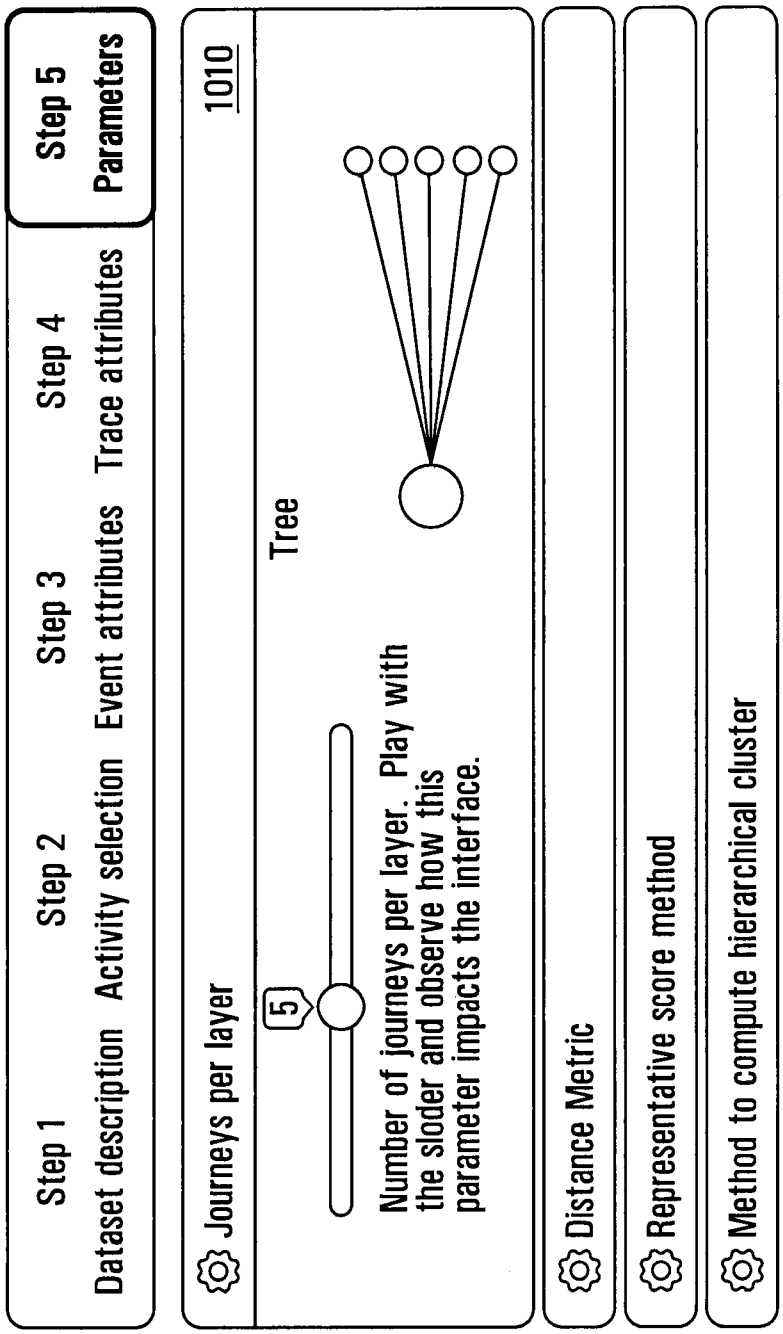
Step 1
Dataset description
Step 2
Activity selection
Step 3
Event attributes
Step 4
Trace attributes
Step 5
Parameters
Journeys per layer
1010
Tree
5
Number of journeys per layer. Play with the sloder and observe how this parameter impacts the interface.
Distance Metric
Representative score method
Method to compute hierarchical cluster

FIG. 11

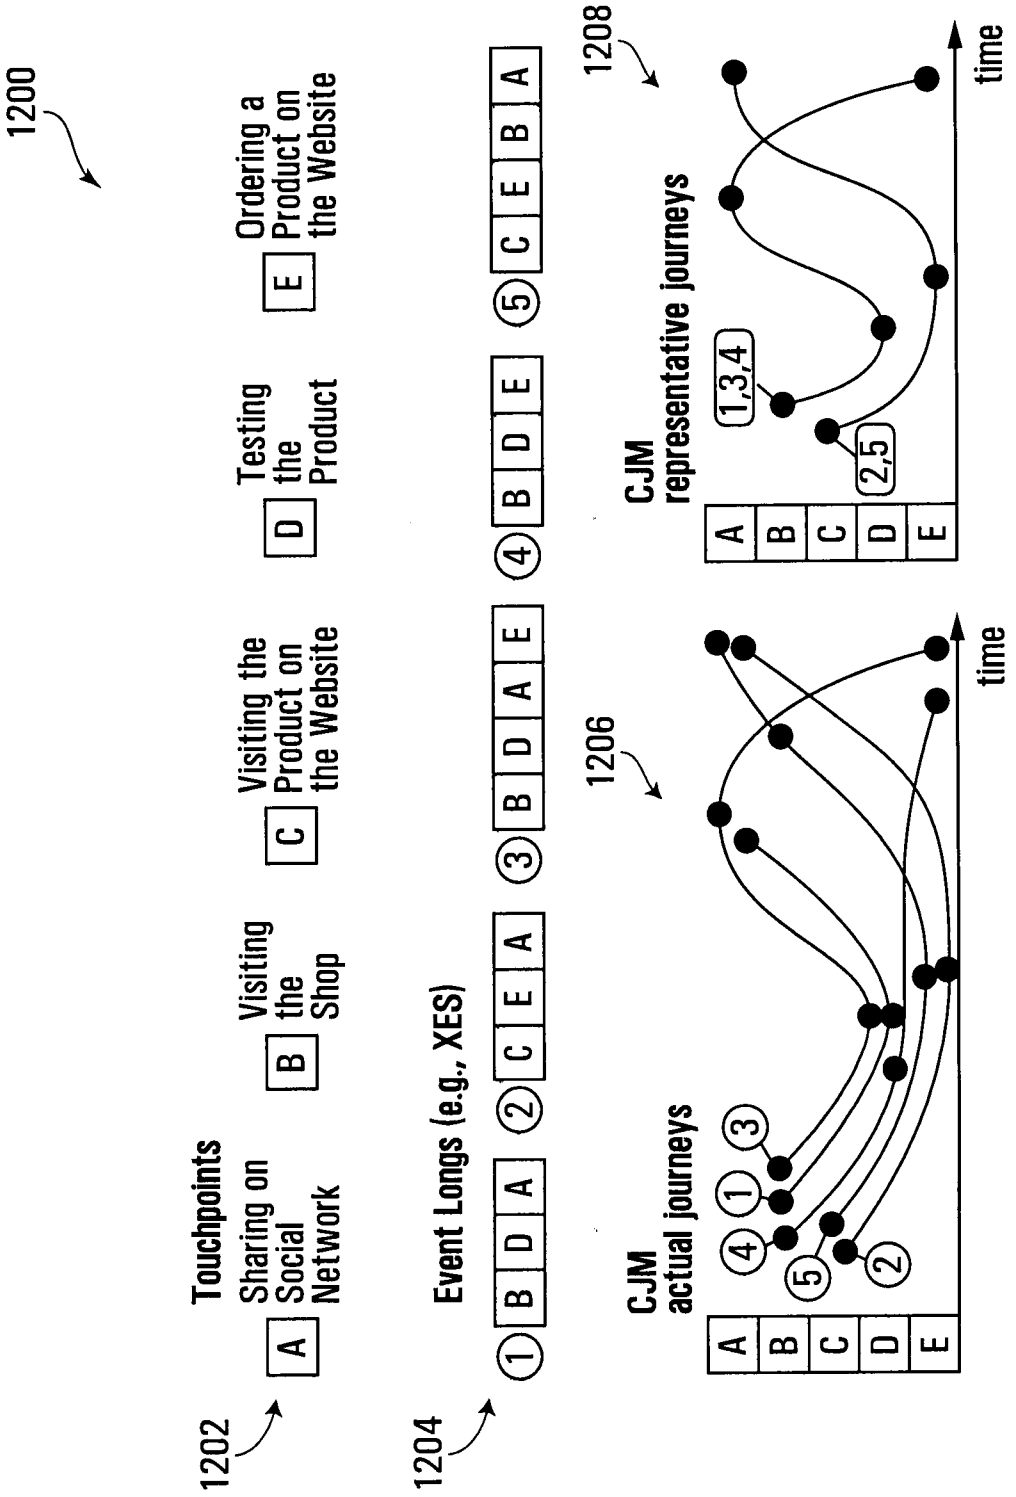
1200
Touchpoints
A
Sharing on Social Network
B
Visiting the Shop
C
Visiting the Product on the Website
D
Testing the Product
E
Ordering a Product on the Website
1202
1204
Event Longs (e.g., XES)
1 B D A
2 C E A
3 B D A E
4 B D E
5 C E B A
1206
CJM actual journeys
1208
CJM representative journeys
1,3,4
2,5
time

FIG. 12

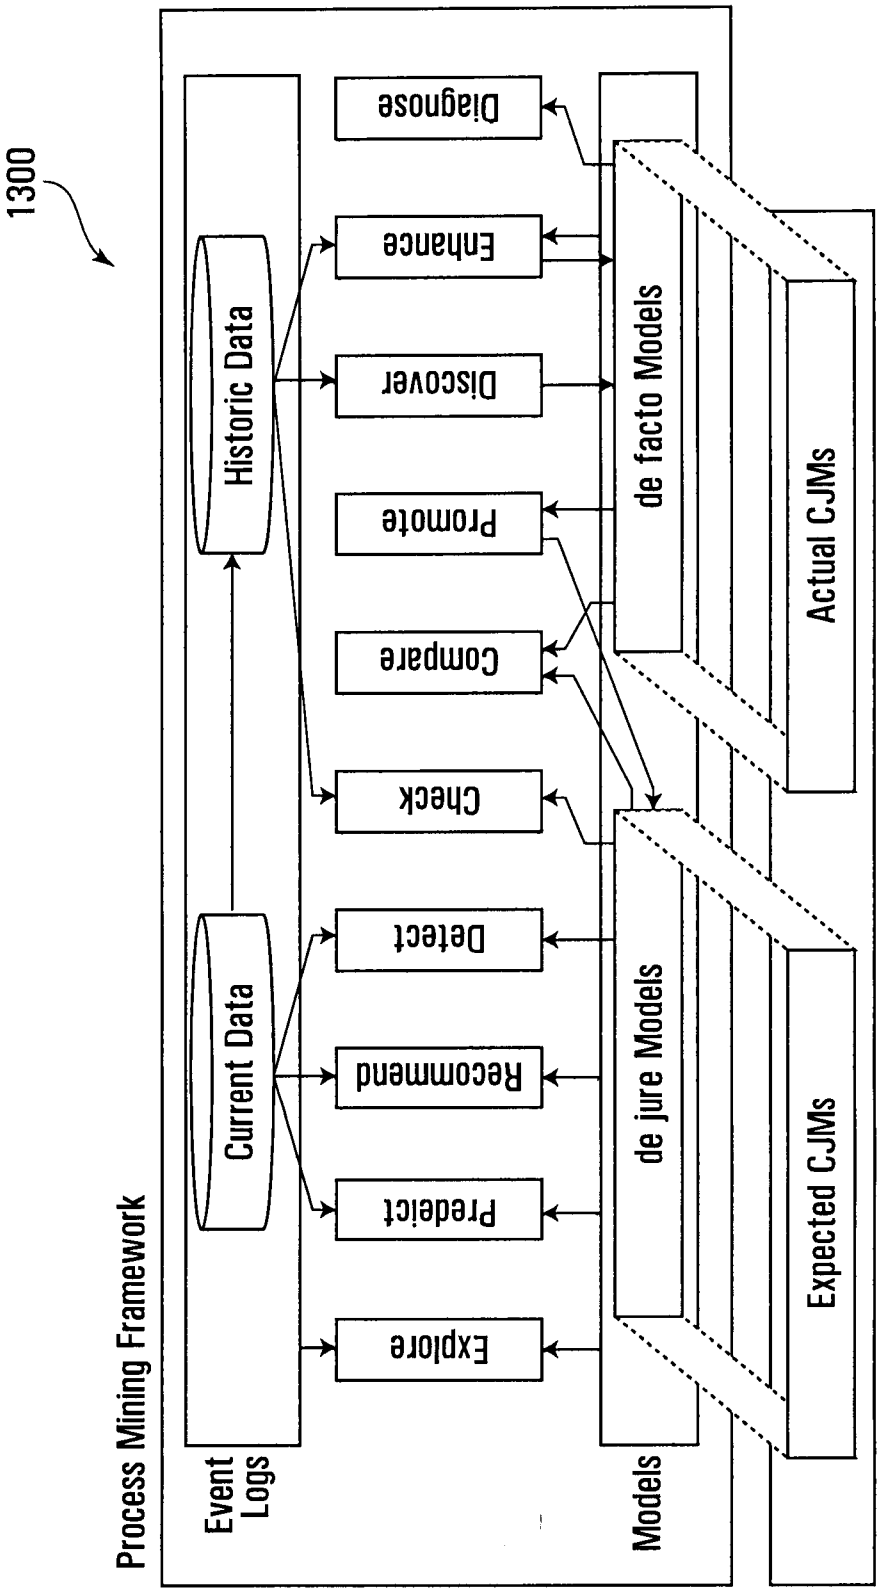
1300
Process Mining Framework
Event Logs
Historic Data
Current Data
Diagnose
Enhance
Discover
Promote
Compare
Check
Detect
Recommend
Predict
Explore
Models
de facto Models
de jure Models
Actual CJMs
Expected CJMs
Customer Journey Mapping Extension

FIG. 13

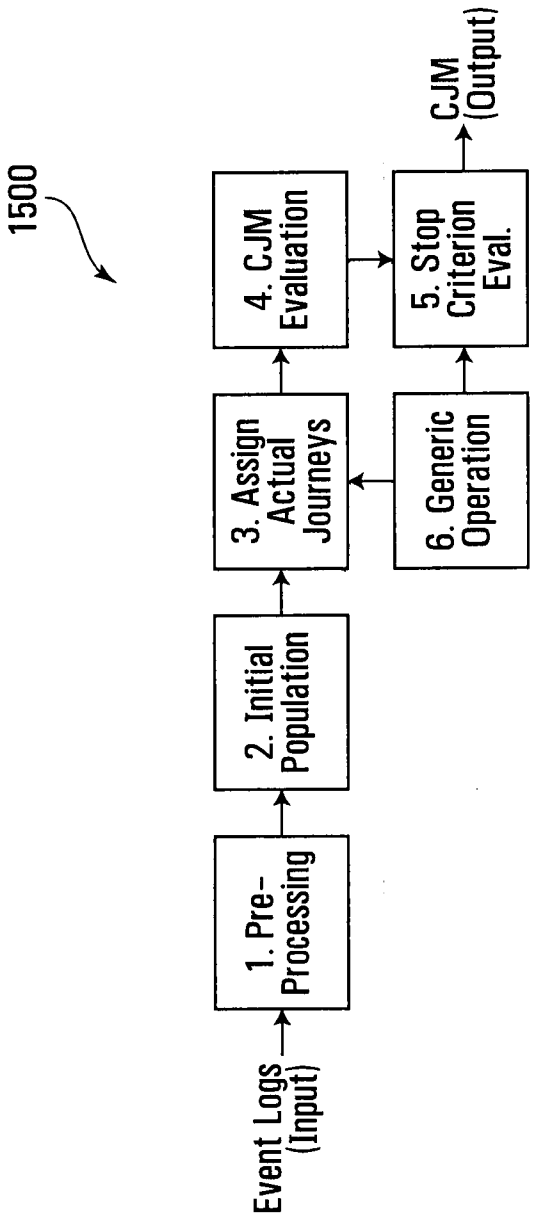
1500
Event Logs (Input)
1. Pre-Processing
2. Initial Population
3. Assign Actual Journeys
4. CJM Evaluation
5. Stop Criterion Eval.
CJM (Output)
6. Generic Operation

FIG. 15

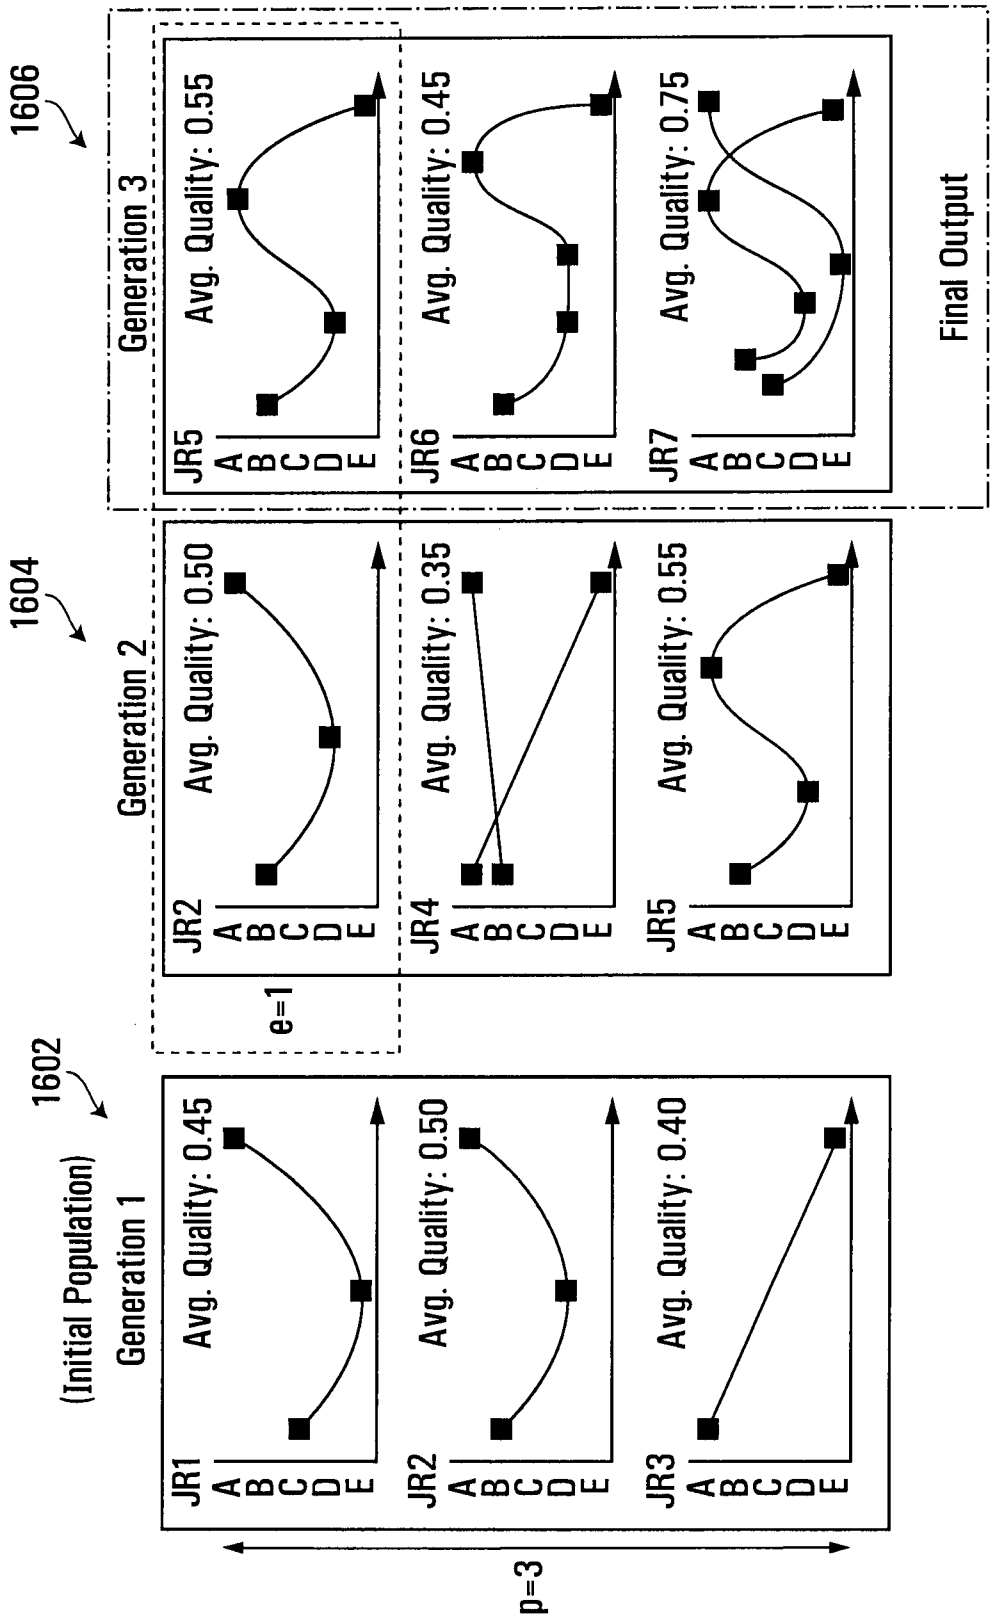
1606
Generation 3
Final Output
JR5 Avg. Quality: 0.55
JR6 Avg. Quality: 0.45
JR7 Avg. Quality: 0.75
1604
Generation 2
e=1
JR2 Avg. Quality: 0.50
JR4 Avg. Quality: 0.35
JR5 Avg. Quality: 0.55
1602
(Initial Population)
Generation 1
JR1 Avg. Quality: 0.45
JR2 Avg. Quality: 0.50
JR3 Avg. Quality: 0.40
A
B
C
D
E
p=3

FIG. 16

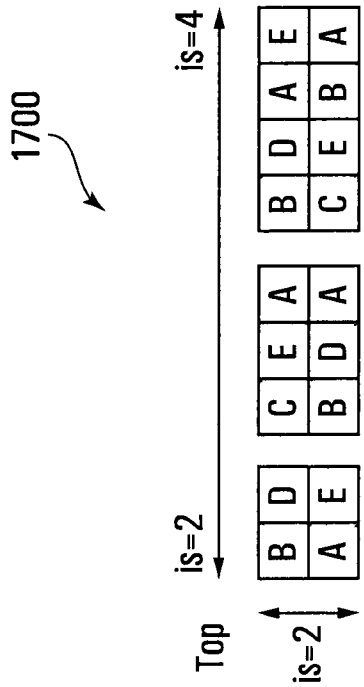

FIG. 17

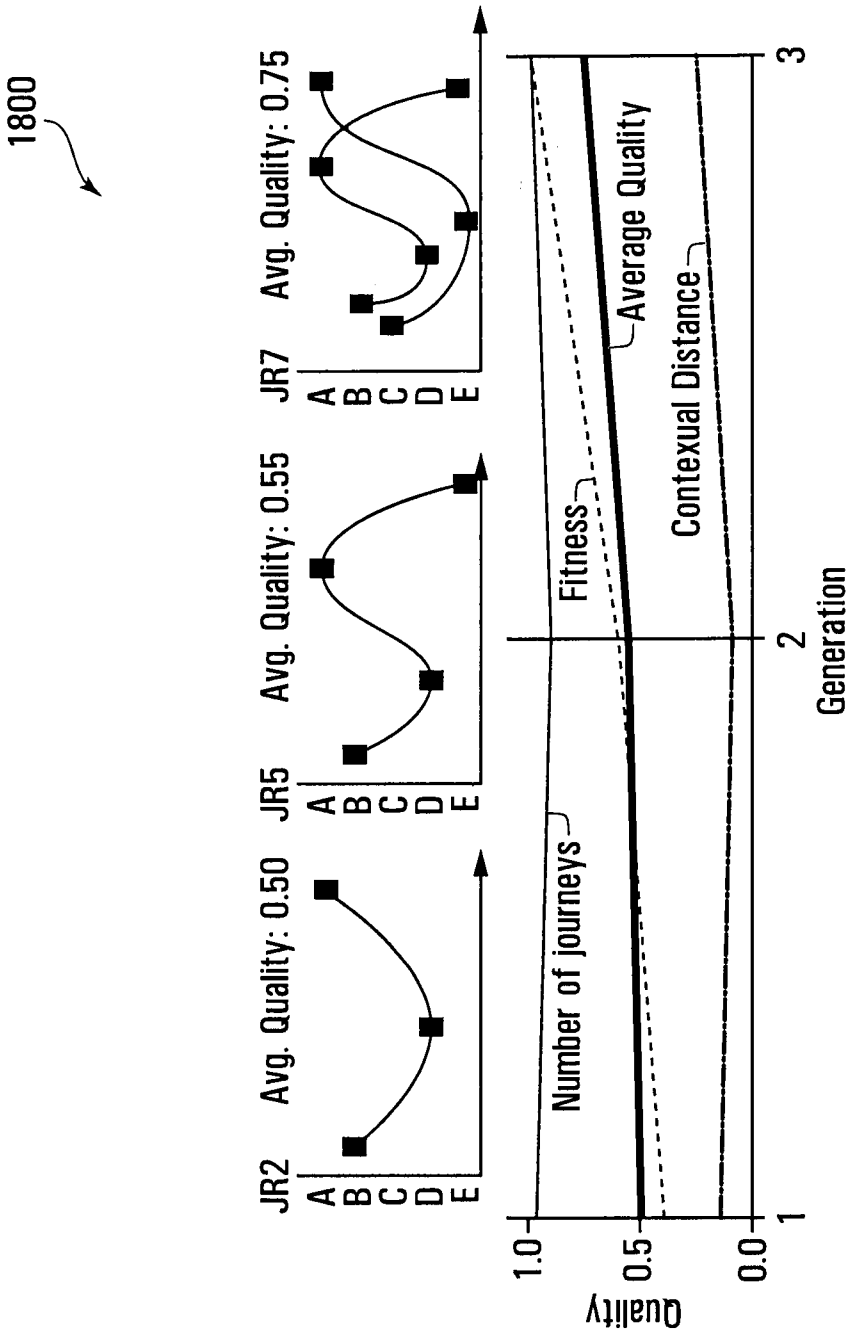
1800
JR2
Avg. Quality: 0.50
JR5
Avg. Quality: 0.55
JR7
Avg. Quality: 0.75
A
B
C
D
E
Quality
1.0
0.5
0.0
1
2
3
Generation
Number of journeys
Fitness
Average Quality
Contexual Distance

FIG. 18

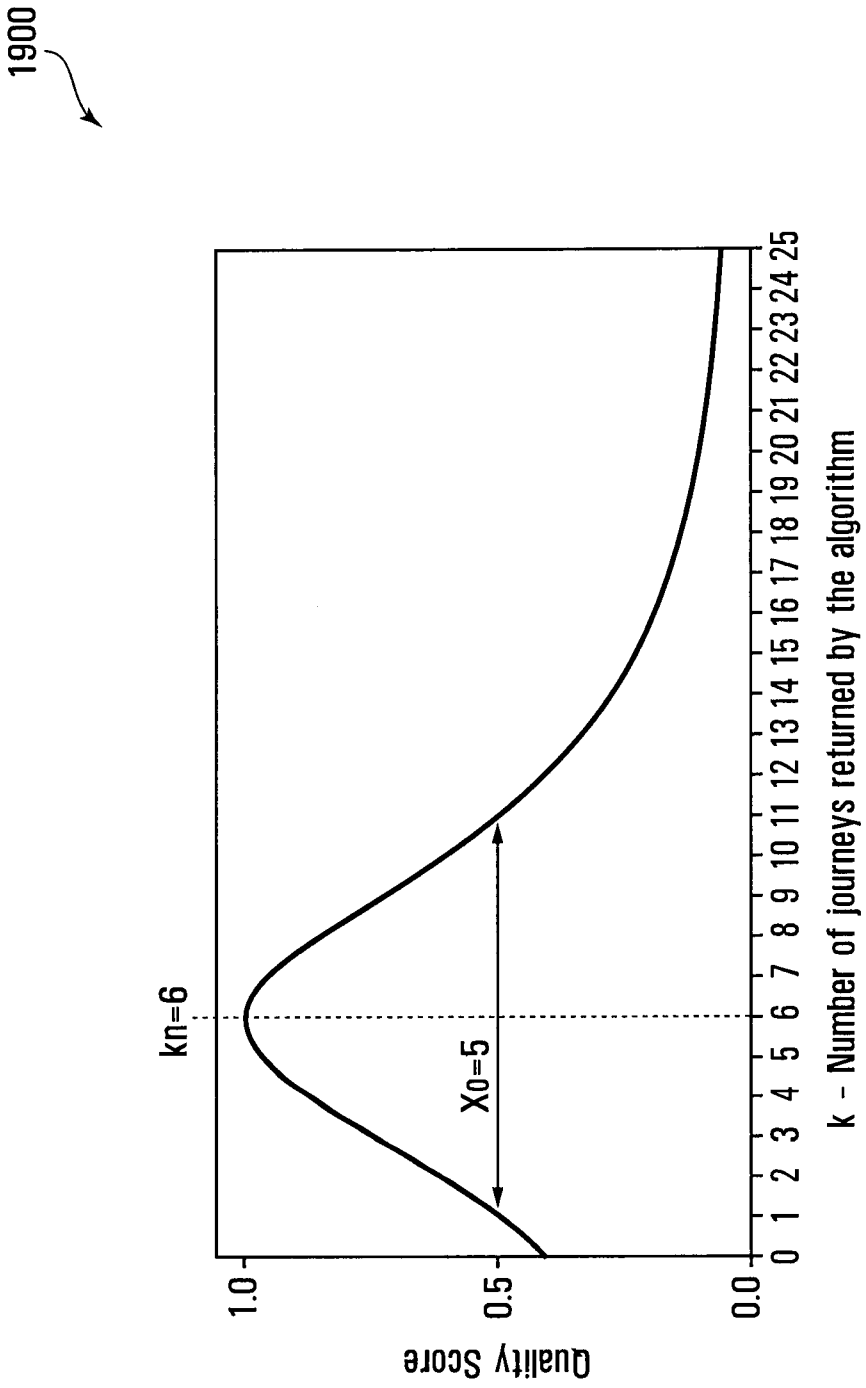
1900
Quality Score
1.0
0.5
0.0
kn=6
X0=5
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25
k - Number of journeys returned by the algorithm

FIG. 19

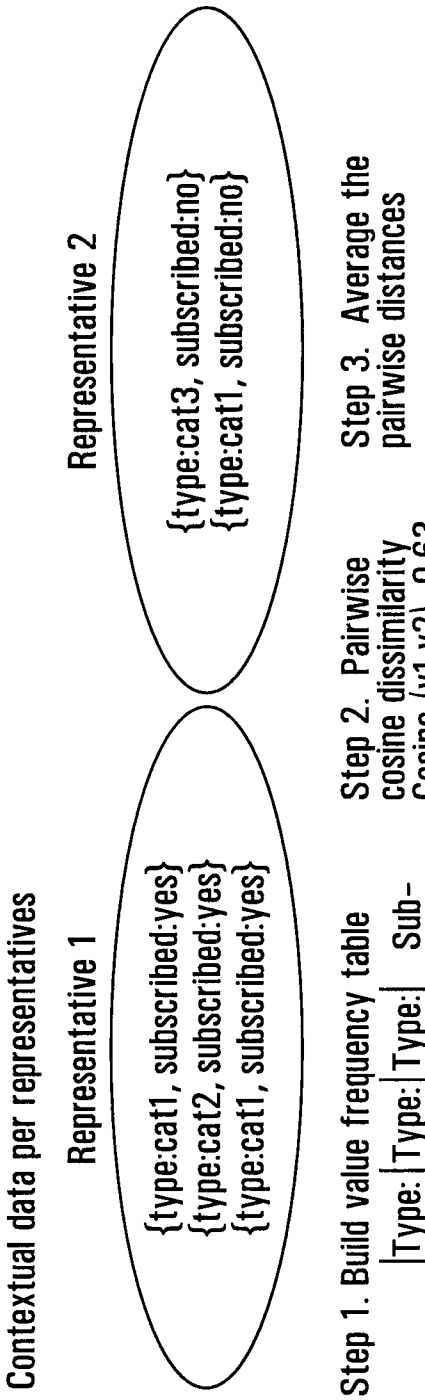
Contextual data per representatives
Representative 1
{type:cat1, subscribed:yes}
{type:cat2, subscribed:yes}
{type:cat1, subscribed:yes}
Representative 2
{type:cat3, subscribed:no}
{type:cat1, subscribed:no}
Step 1. Build value frequency table
Type: cat1
Type: cat2
Type: cat3
Sub-scribed?
Rep1: v1 2 1 0 0
Rep2: v2 1 0 1 2
Step 2. Pairwise cosine dissimilarity
Cosine (v1,v2)=0.63
Step 3. Average the pairwise distances

FIG. 20

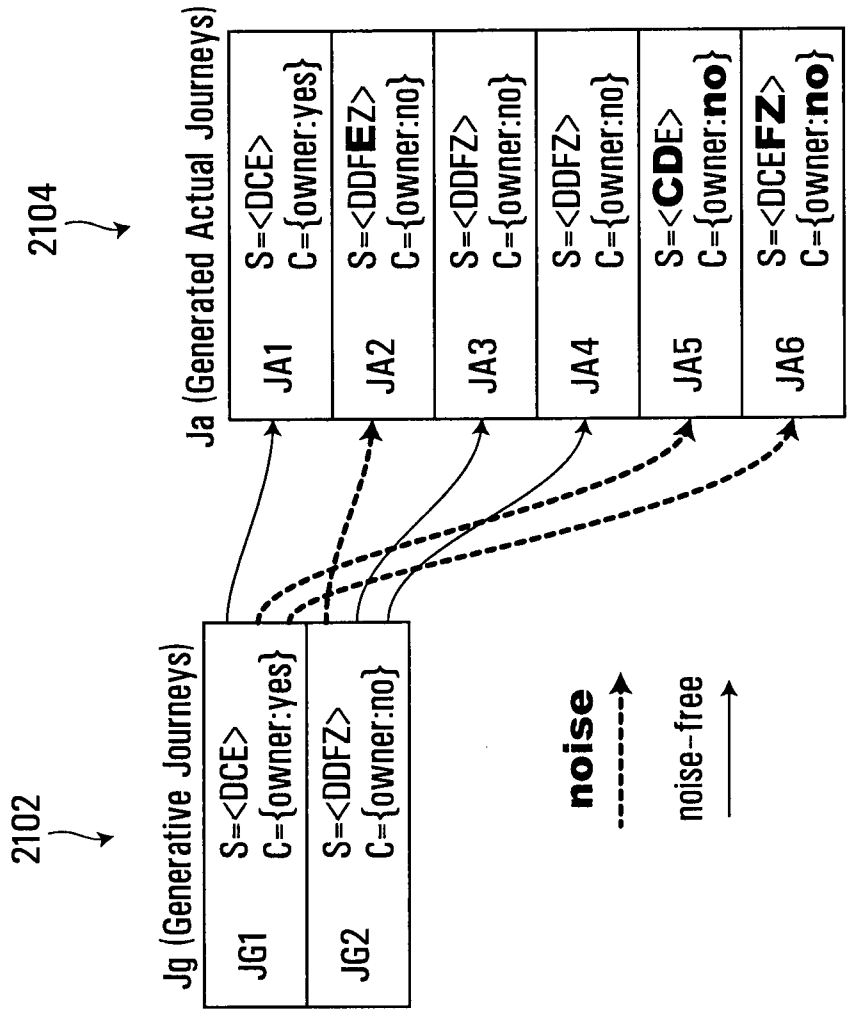
2102
Jg (Generative Journeys)
JG1 S=<DCE> C={owner:yes}
JG2 S=<DDFZ> C={owner:no}
2104
Ja (Generated Actual Journeys)
JA1 S=<DCE> C={owner:yes}
JA2 S=<DDFEZ> C={owner:no}
JA3 S=<DDFZ> C={owner:no}
JA4 S=<DDFZ> C={owner:no}
JA5 S=<CDE> C={owner:no}
JA6 S=<DCEFZ> C={owner:no}
noise
noise-free

FIG. 21

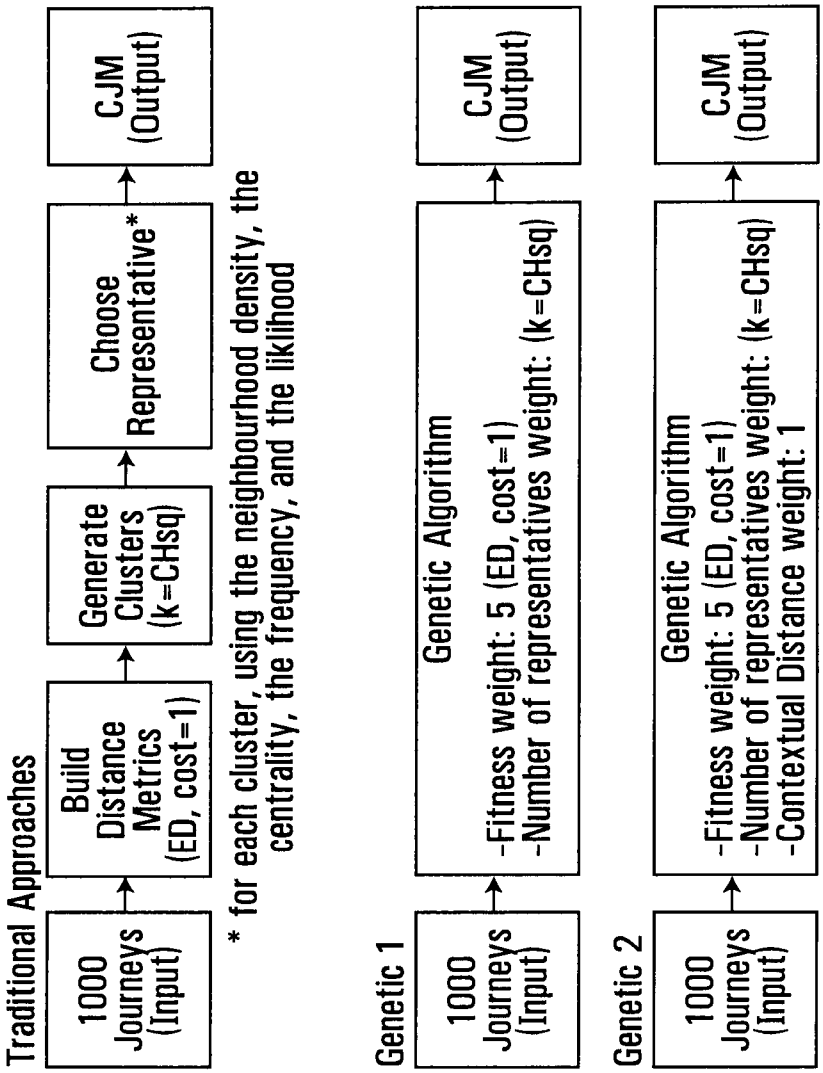
Traditional Approaches
1000 Journeys (Input)
Build Distance Metrics (ED, cost=1)
Generate Clusters (k=CHsq)
Choose Representative*
CJM (Output)
* for each cluster, using the neighbourhood density, the centrality, the frequency, and the likilihood
Genetic 1
1000 Journeys (Input)
Genetic Algorithm
-Fitness weight: 5 (ED, cost=1)
-Number of representatives weight: (k=CHsq)
CJM (Output)
Genetic 2
1000 Journeys (Input)
Genetic Algorithm
-Fitness weight: 5 (ED, cost=1)
-Number of representatives weight: (k=CHsq)
-Contextual Distance weight: 1
CJM (Output)

FIG. 22

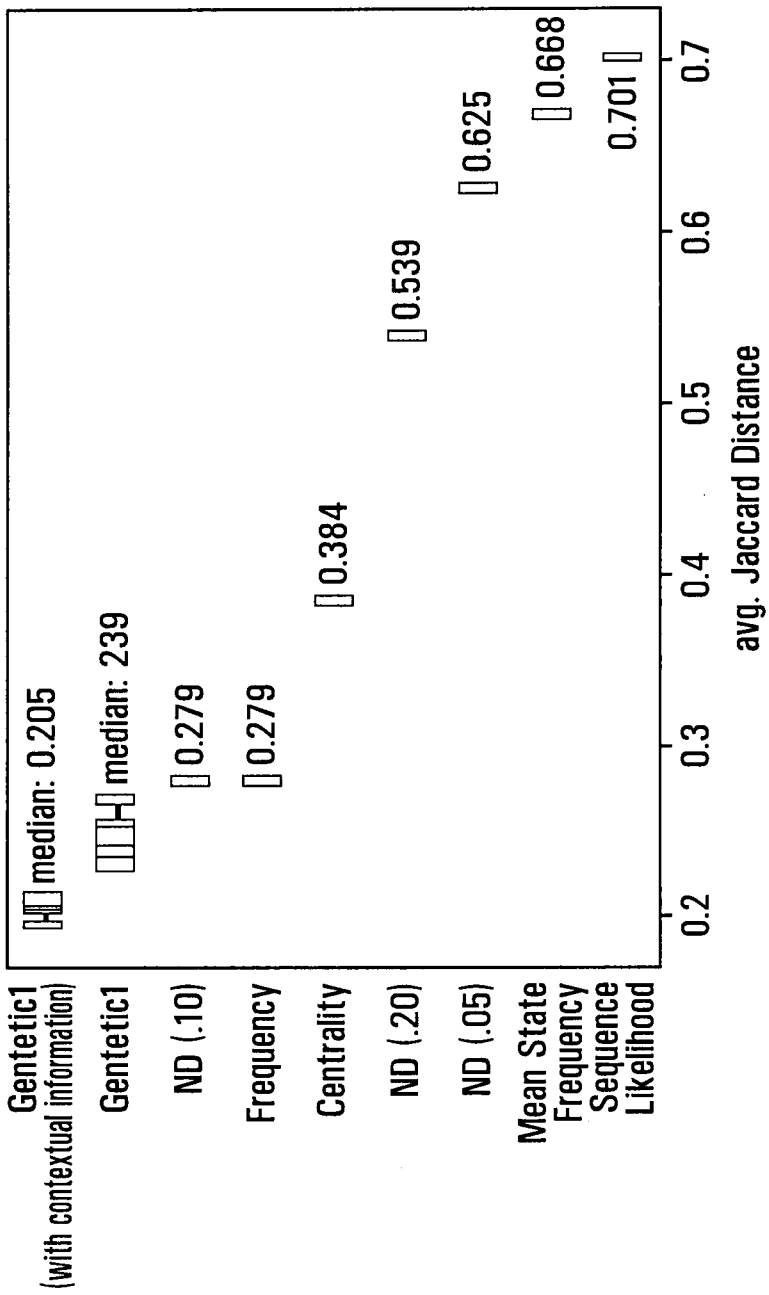
Gentetic1 (with contextual information)
median: 0.205
Gentetic1
median: 239
ND (.10)
0.279
Frequency
0.279
Centrality
0.384
ND (.20)
0.539
ND (.05)
0.625
Mean State Frequency Sequence Likelihood
0.668
0.701
0.2
0.3
0.4
0.5
0.6
0.7
avg. Jaccard Distance

FIG. 23

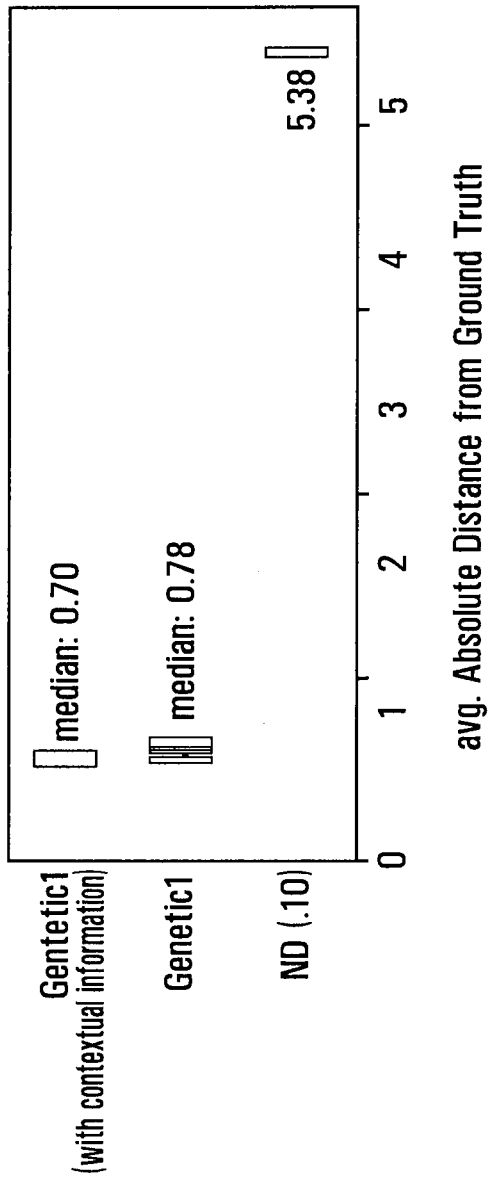
Gentetic1
(with contextual information)
Genetic1
ND (.10)
median: 0.70
median: 0.78
5.38
0
1
2
3
4
5
avg. Absolute Distance from Ground Truth

FIG. 24

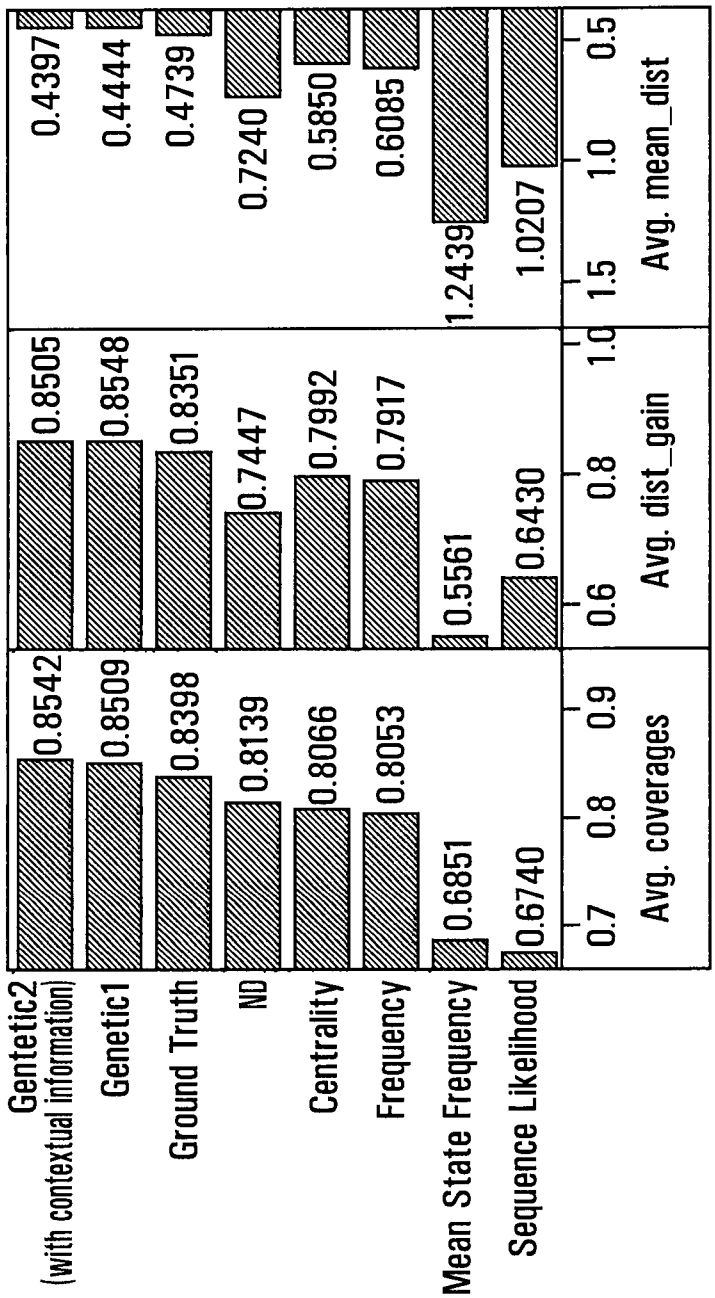
Gentetic2 (with contextual information)
Genetic1
Ground Truth
ND
Centrality
Frequency
Mean State Frequency
Sequence Likelihood
0.8542
0.8509
0.8398
0.8139
0.8066
0.8053
0.6851
0.6740
0.8505
0.8548
0.8351
0.7447
0.7992
0.7917
0.5561
0.6430
0.4397
0.4444
0.4739
0.7240
0.5850
0.6085
1.2439
1.0207
0.7 0.8 0.9
Avg. coverages
0.6 0.8 1.0
Avg. dist_gain
1.5 1.0 0.5
Avg. mean_dist

FIG. 25

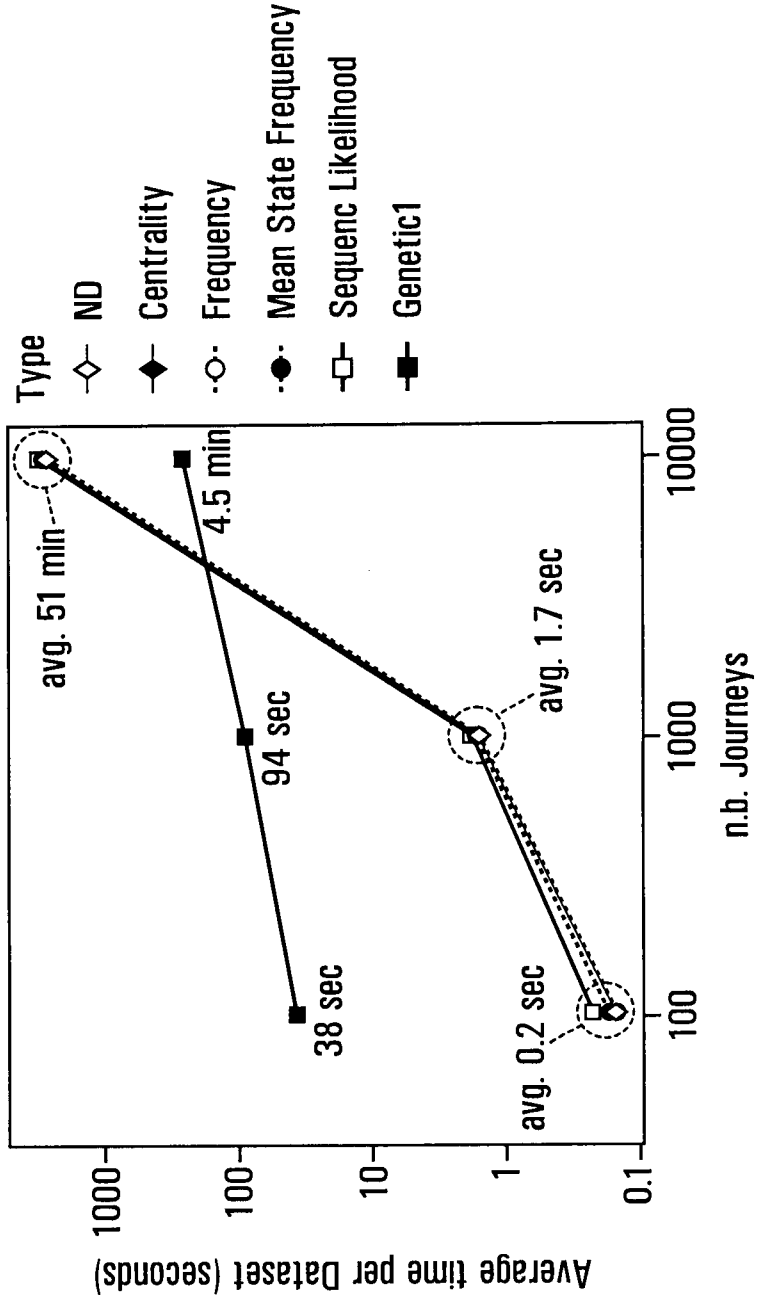
Type
ND
Centrality
Frequency
Mean State Frequency
Sequenc Likelihood
Genetic1
avg. 51 min
4.5 min
94 sec
38 sec
avg. 1.7 sec
avg. 0.2 sec
1000
100
10
1
0.1
100
1000
10000
n.b. Journeys
Average time per Dataset (seconds)

FIG. 26

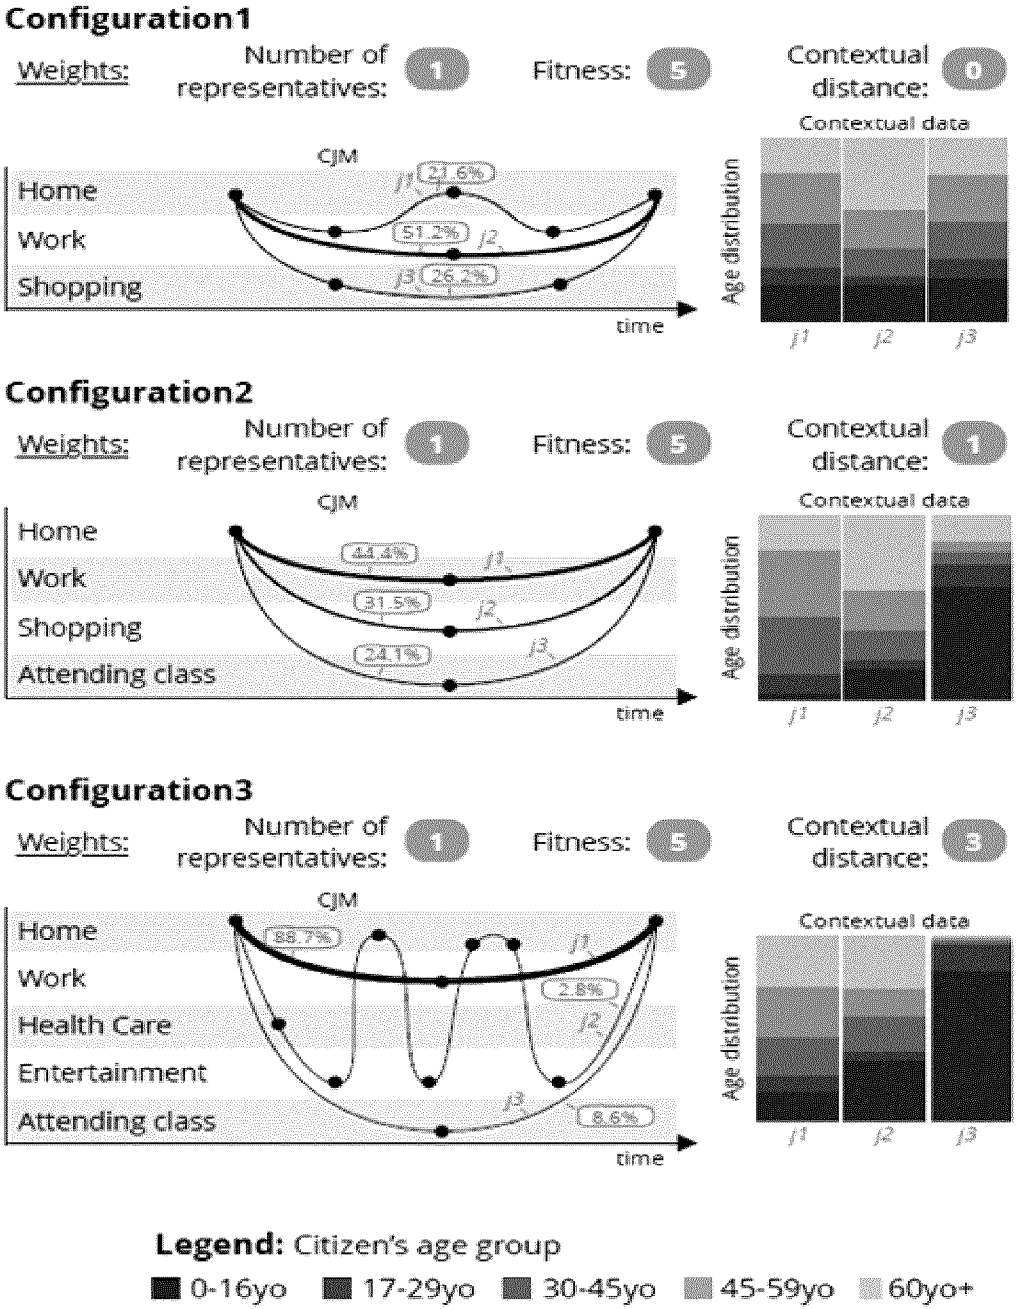
Configuration1
Weights:
Number of representatives: 1
Fitness: 5
Contextual distance: 0
CJM
Home
Work
Shopping
21.6%
51.2%
26.2%
time
Contextual data
Age distribution
Configuration2
Weights:
Number of representatives: 1
Fitness: 5
Contextual distance: 1
CJM
Home
Work
Shopping
Attending class
44.4%
31.5%
24.1%
time
Contextual data
Age distribution
Configuration3
Weights:
Number of representatives: 1
Fitness: 5
Contextual distance: 3
CJM
Home
Work
Health Care
Entertainment
Attending class
88.7%
2.8%
8.6%
time
Contextual data
Age distribution
Legend: Citizen's age group
0-16yo
17-29yo
30-45yo
45-59yo
60yo+

FIG. 27

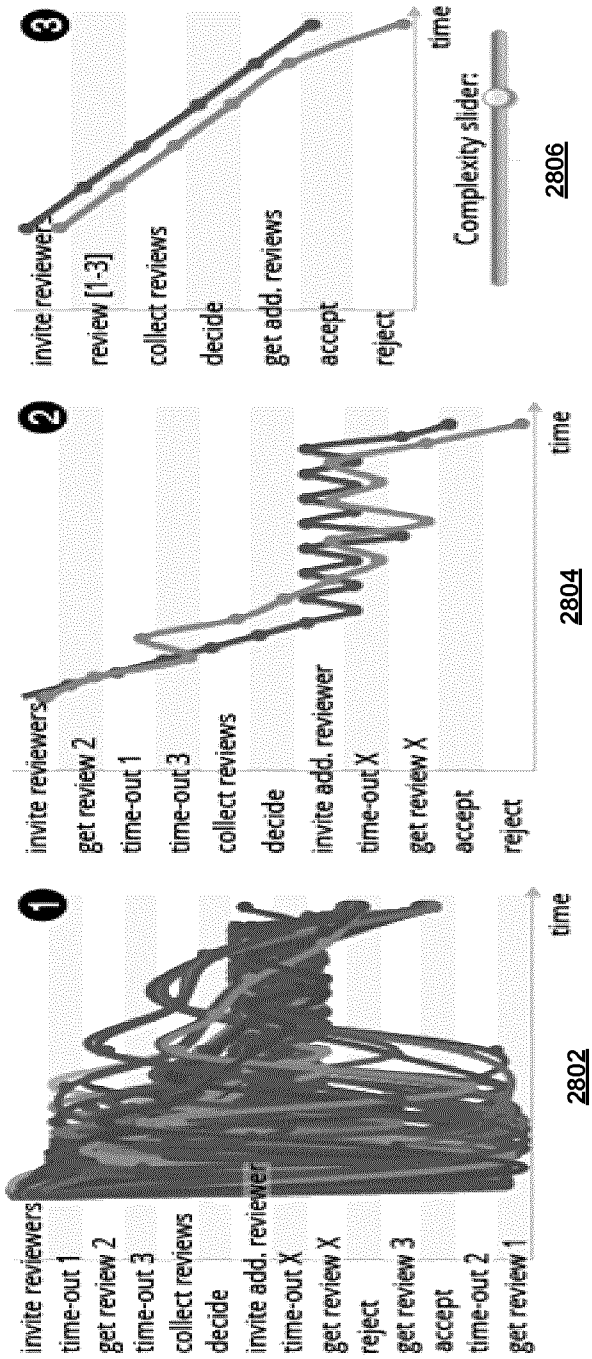
1
invite reviewers
time-out 1
get review 2
time-out 3
collect reviews
decide
invite add. reviewer
time-out X
get review X
reject
get review 3
accept
time-out 2
get review 1
time
2802
2
invite reviewers
get review 2
time-out 1
time-out 3
collect reviews
decide
invite add. reviewer
time-out X
get review X
accept
reject
time
2804
3
invite reviewers
review [1-3]
collect reviews
decide
get add. reviews
accept
reject
time
Complexity slider:
2806

FIG. 28

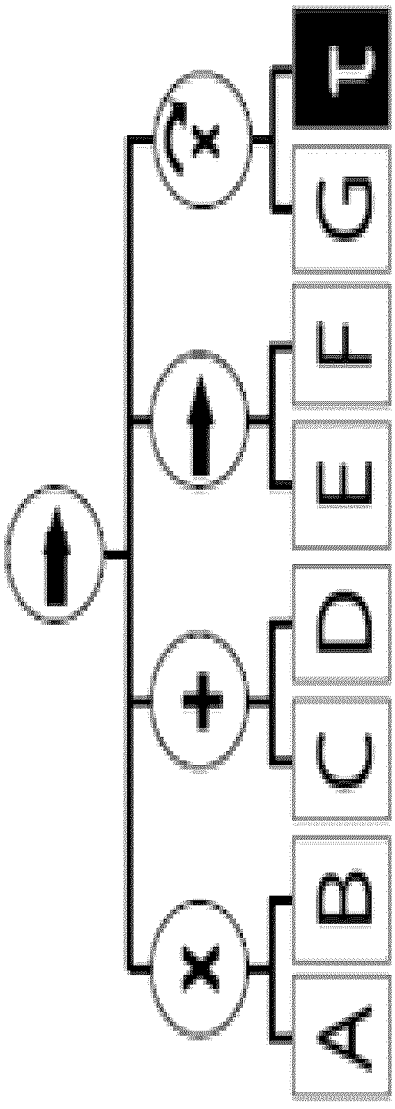

FIG. 29

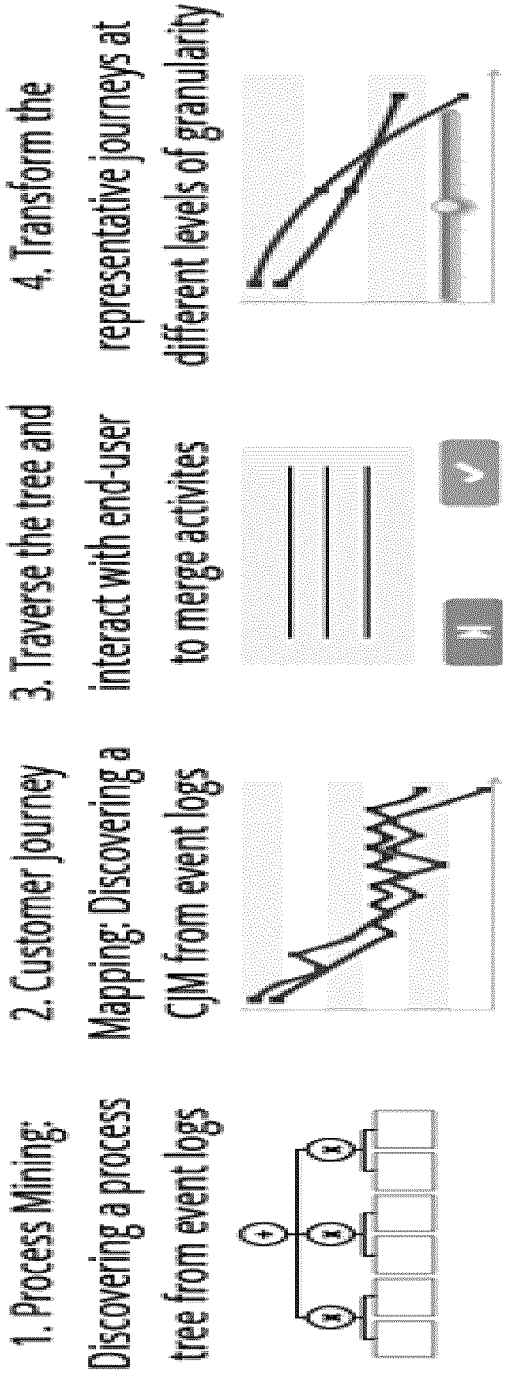
1. Process Mining: Discovering a process tree from event logs
2. Customer Journey Mapping: Discovering a CJM from event logs
3. Traverse the tree and interact with end-user to merge activites
4. Transform the representative journeys at different levels of granularity

3000

FIG. 30

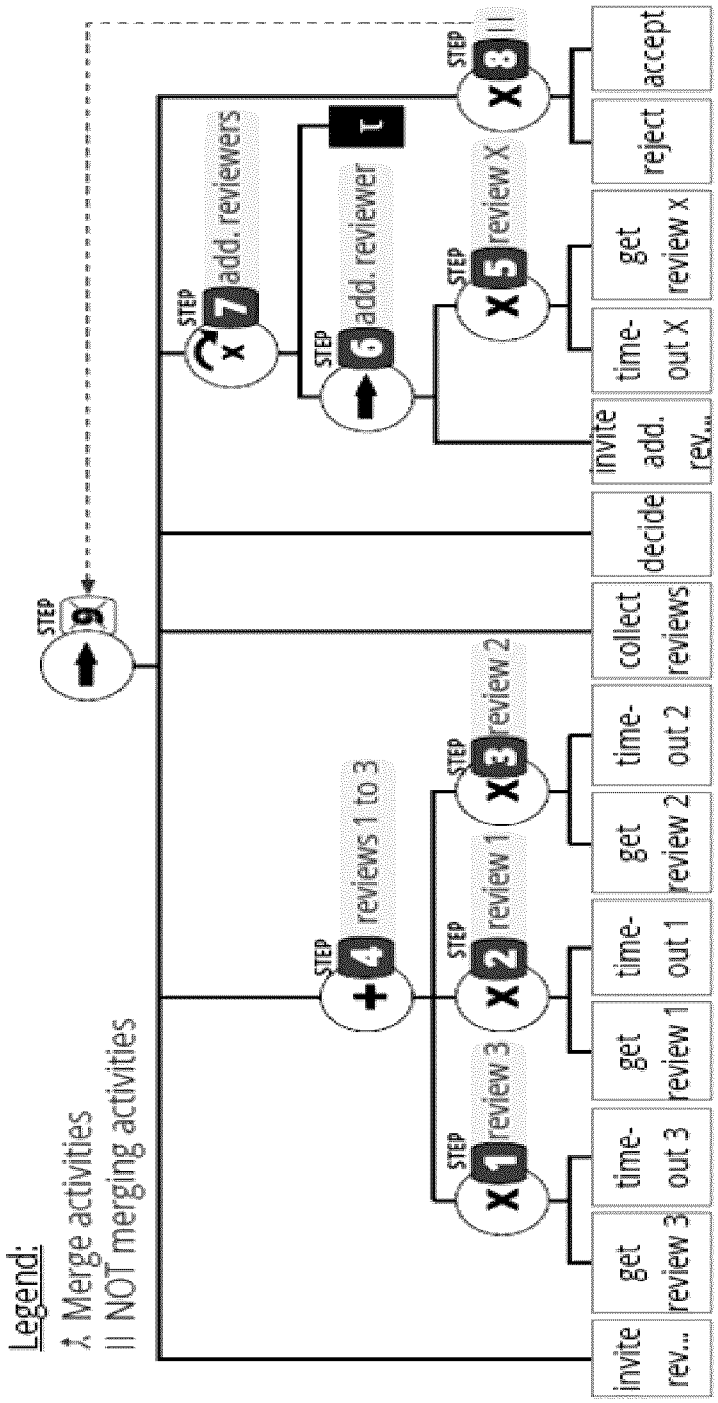
Legend:
∧ Merge activities
|| NOT merging activities
STEP 9
STEP 4 reviews 1 to 3
STEP 1 review 3
STEP 2 review 1
STEP 3 review 2
STEP 7 add. reviewers
STEP 6 add. reviewer
STEP 5 review X
STEP 8 ||
invite rev...
get review 3
time-out 3
get review 1
time-out 1
get review 2
time-out 2
collect reviews
decide
invite add. rev...
time-out X
get review x
reject
accept
1

3100

FIG. 31

Make CJM simpler with process tree
According to the process tree, the following activities are in a xor relation:
• get review 3
• time-out 3
Does it make sense to merge them into a single activity ?
No
review 3
Yes
3200

FIG. 32

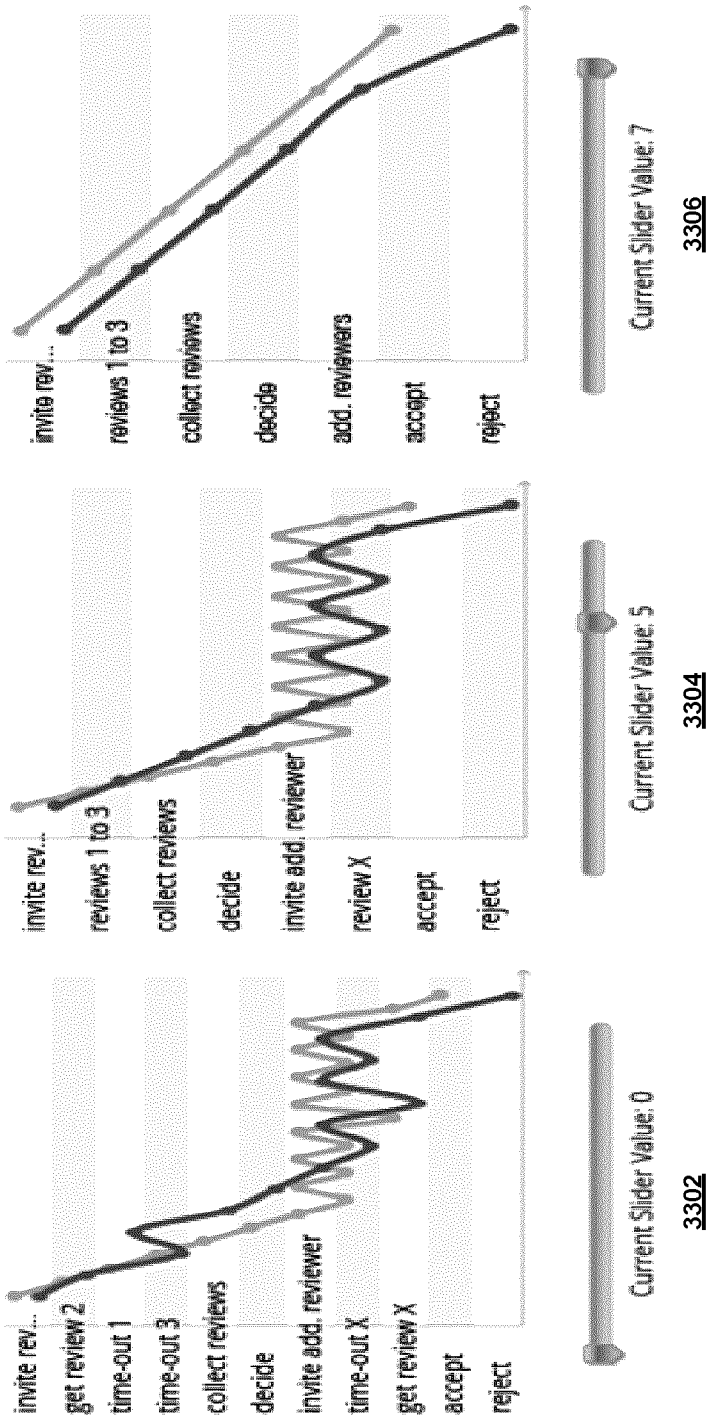
invite rev...
get review 2
time-out 1
time-out 3
collect reviews
decide
invite add. reviewer
time-out X
get review X
accept
reject
Current Slider Value: 0
3302
invite rev...
reviews 1 to 3
collect reviews
decide
invite add. reviewer
review X
accept
reject
Current Slider Value: 5
3304
invite rev...
reviews 1 to 3
collect reviews
decide
add. reviewers
accept
reject
Current Slider Value: 7
3306

FIG. 33

INTERNATIONAL SEARCH REPORT

International application No.
PCT/CA2019/050823

A. CLASSIFICATION OF SUBJECT MATTER
IPC: ***G06F 16/906*** (2019.01) , ***G06F 16/80*** (2019.01) , ***G06F 16/904*** (2019.01) , ***G06N 20/00*** (2019.01) , *G16H 50/70* (2018.01) , *G06N 3/02* (2006.01) (more IPCs on the last page)

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC: *G06F 16/906* (2019.01) , *G06F 16/80* (2019.01) , *G06F 16/904* (2019.01) , *G06N 20/00* (2019.01) , *G16H 50/70* (2018.01) , *G06N 3/02* (2006.01) , *G06N 3/12* (2006.01)

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic database(s) consulted during the international search (name of database(s) and, where practicable, search terms used)
Questel Orbit™, Canadian Patent Database, IEEEXplore™, Google Scholar™, Google Patents™ and the search strings: cluster, process mining, journey map, process discovery, business process management, data mining, event, log, event trace, event clustering, trace clustering, tree, hierarchy, process model, visualization, display, representation, gui

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US20170111245(ISHAKIAN et al.) 20 April 2017(20-04-2017) *abstract; paragraphs 0009-0022, 0081-0082; fig. 4*	1 – 41
A	US9275359(FOMIN et al.) 1 March 2016(01-03-2016) *abstract; col. 2, line 6-45; col. 4, lines 20-41; col. 5, line 6- col. 6, line 2; 0009-0022, 0081-0082; fig. 3*	1 – 41
A	US20140058789(DOEHRING et al.) 27 February 2014(27-02-2014) *abstract; paragraphs 0004, 0016-0023; fig. 1*	1 – 41
A	US9324038(CHARIF et al.) 26 April 2016(26-04-2016) *abstract; col. 1, line 45-col. 4, line 29; figs. 1, 2*	1 – 41

☐ Further documents are listed in the continuation of Box C.

☑ See patent family annex.

*	Special categories of cited documents:
"A"	document defining the general state of the art which is not considered to be of particular relevance
"D"	document cited by the applicant in the international application
"E"	earlier application or patent but published on or after the international filing date
"L"	document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O"	document referring to an oral disclosure, use, exhibition or other means
"P"	document published prior to the international filing date but later than the priority date claimed
"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"&"	document member of the same patent family

Date of the actual completion of the international search 15 August 2019 (15-08-2019)	Date of mailing of the international search report 15 August 2019 (15-08-2019)
Name and mailing address of the ISA/CA Canadian Intellectual Property Office Place du Portage I, C114 - 1st Floor, Box PCT 50 Victoria Street Gatineau, Quebec K1A 0C9 Facsimile No.: 819-953-2476	Authorized officer Dan Marinescu (819) 639-8202

Box No. II **Observations where certain claims were found unsearchable (Continuation of item 2 of the first sheet)**

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claim Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claim Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claim Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III **Observations where unity of invention is lacking (Continuation of item 3 of first sheet)**

This International Searching Authority found multiple inventions in this international application, as follows:

A. Claims 1 – 18 and 22 – 41: computing clusters of event traces, each cluster corresponding to a set of event traces and a representative event trace based on representative attributes of the set of event traces; visualizing the elements of the clusters.
B. Claim 19: generating a mapping of process mining elements to event traces, the mapping linking process models to customer experiences; generating a data structure for hierarchical components of a journey map; visualizing elements for the components of the journey map.
C. Claims 20 – 21: generating a process tree for event traces by grouping the event traces based on similar attributes; parsing the process tree; merging a set of event traces that belong to a subset of the parsed tree; visualizing elements representing an abstract process tree.

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.

2. ☑ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.

3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claim Nos.:

4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claim Nos.:

Remark on Protest

☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.

☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.

☐ No protest accompanied the payment of additional search fees.

Patent Document Cited in Search Report	Publication Date	Patent Family Member(s)	Publication Date
US2017111245A1	20 April 2017 (20-04-2017)	US2017111245A1 US2017286190A1	20 April 2017 (20-04-2017) 05 October 2017 (05-10-2017)
US9275359B1	01 March 2016 (01-03-2016)	None	
US2014058789A1	27 February 2014 (27-02-2014)	None	
US9324038B2	26 April 2016 (26-04-2016)	US2015142707A1 US9324038B2	21 May 2015 (21-05-2015) 26 April 2016 (26-04-2016)

G06N 3/12 (2006.01)