



US 20140379619A1

(19) **United States**(12) **Patent Application Publication**
Permeh et al.(10) **Pub. No.: US 2014/0379619 A1**(43) **Pub. Date: Dec. 25, 2014**(54) **AUTOMATED SYSTEM FOR GENERATIVE
MULTIMODEL MULTICLASS
CLASSIFICATION AND SIMILARITY
ANALYSIS USING MACHINE LEARNING****Publication Classification**

(51) **Int. Cl.**
G06N 99/00 (2006.01)
G06N 5/02 (2006.01)

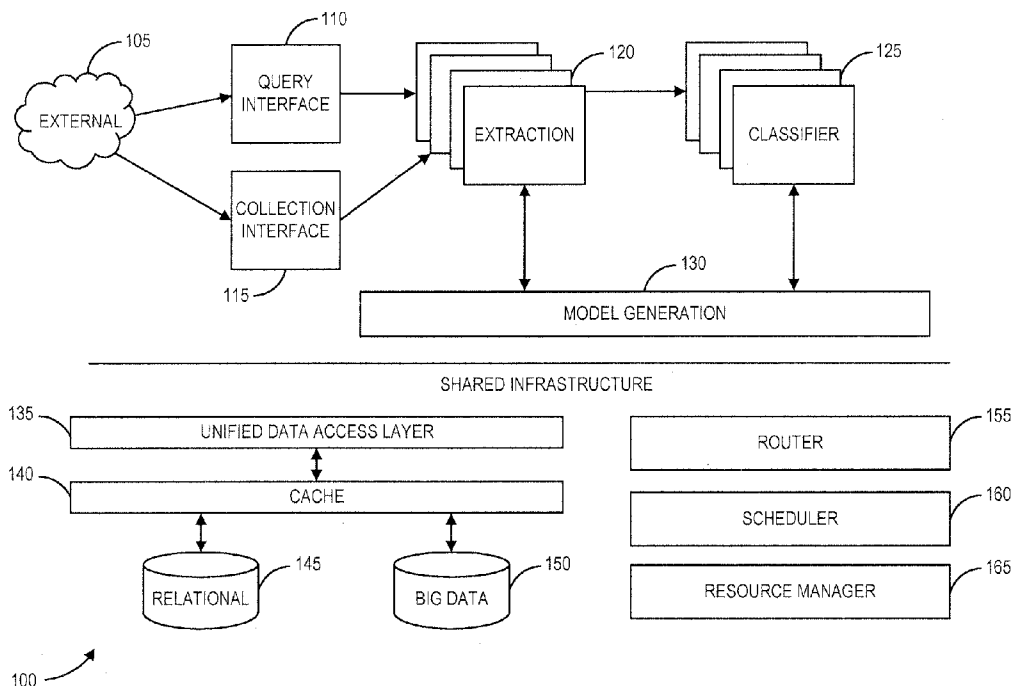
(52) **U.S. Cl.**
CPC **G06N 99/005** (2013.01); **G06N 5/02** (2013.01)
USPC **706/12**

(71) Applicant: **Cylance Inc.**, Irvine, CA (US)(72) Inventors: **Ryan Permeh**, Irvine, CA (US); **Stuart McClure**, Irvine, CA (US); **Matthew Wolff**, Newport Beach, CA (US); **Gary Golomb**, Santa Cruz, CA (US); **Derek A. Soeder**, Irvine, CA (US); **Seagen Levites**, Portland, OR (US); **Michael O'Dea**, Estero, FL (US); **Gabriel Acevedo**, Irvine, CA (US); **Glenn Chisholm**, Irvine, CA (US)(21) Appl. No.: **14/313,863**(22) Filed: **Jun. 24, 2014****Related U.S. Application Data**

(60) Provisional application No. 61/838,820, filed on Jun. 24, 2013.

(57) **ABSTRACT**

A sample of data is placed within a directed graph that comprises a plurality of hierarchical nodes that form a queue of work items for a particular worker class that are used to process the sample of data. Subsequently, work items are scheduled within the queue for each of a plurality of workers by traversing the nodes of the directed graph. The work items are then served to the workers according to the queue. Results can later be received from the workers for the work items (the nodes of the directed graph are traversed based on the received results). In addition, in some variations, the results can be classified so that one or models can be generated. Related systems, methods, and computer program products are also described.






FIG. 1

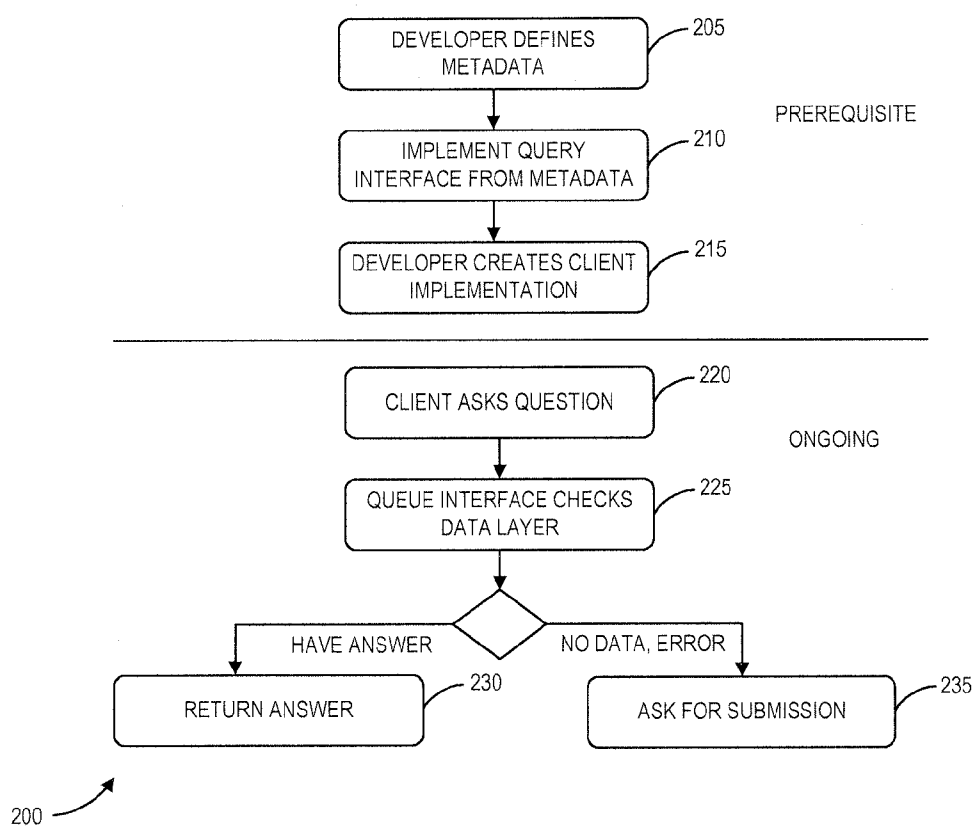


FIG. 2




FIG. 3

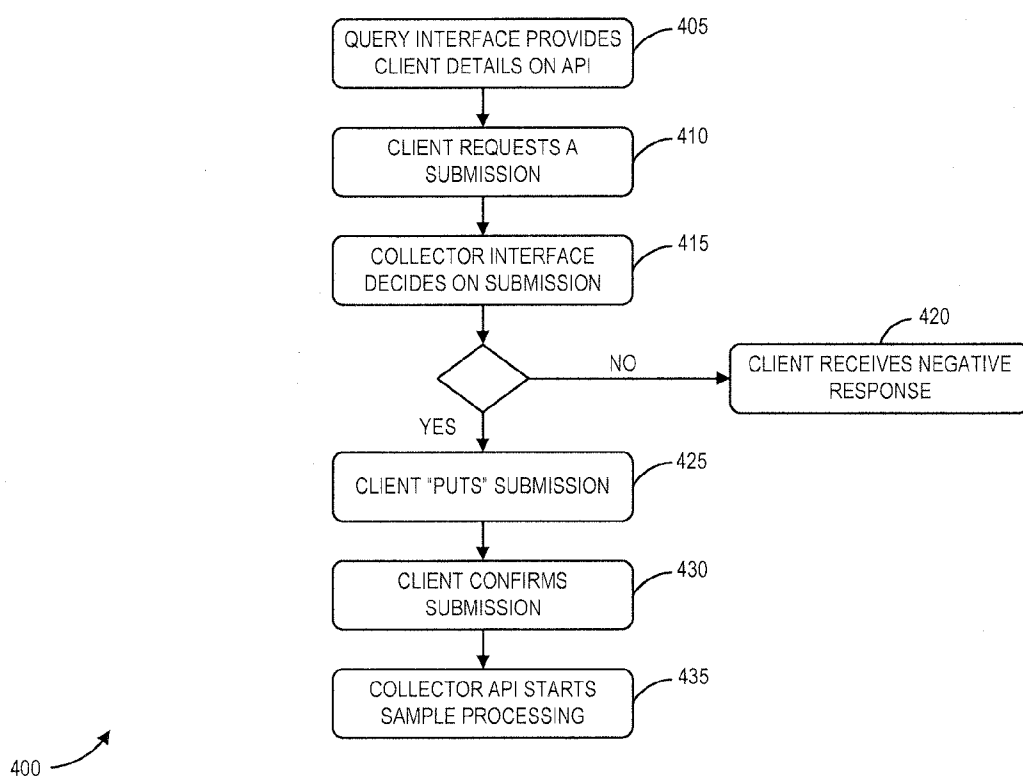
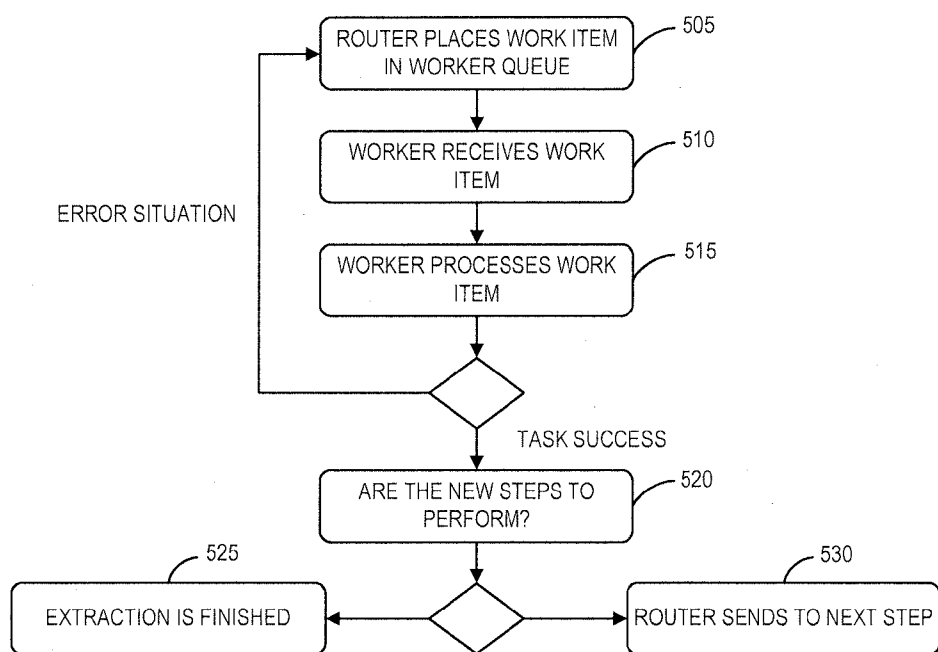


FIG. 4



500 ↗

FIG. 5

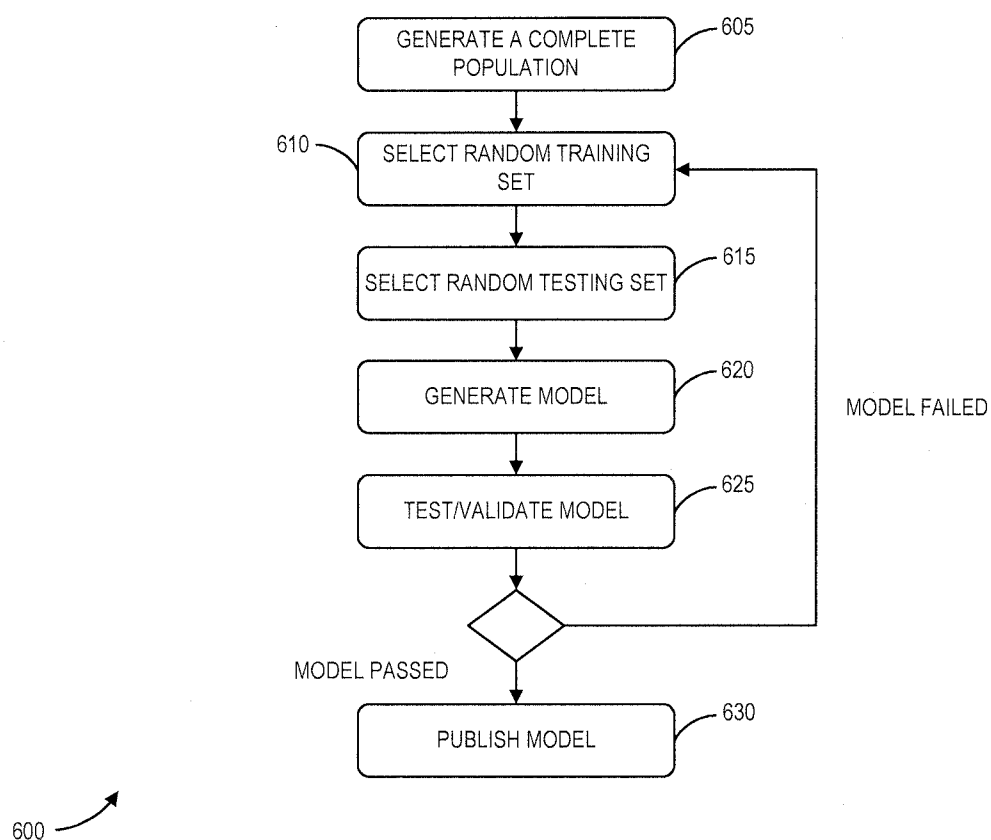


FIG. 6

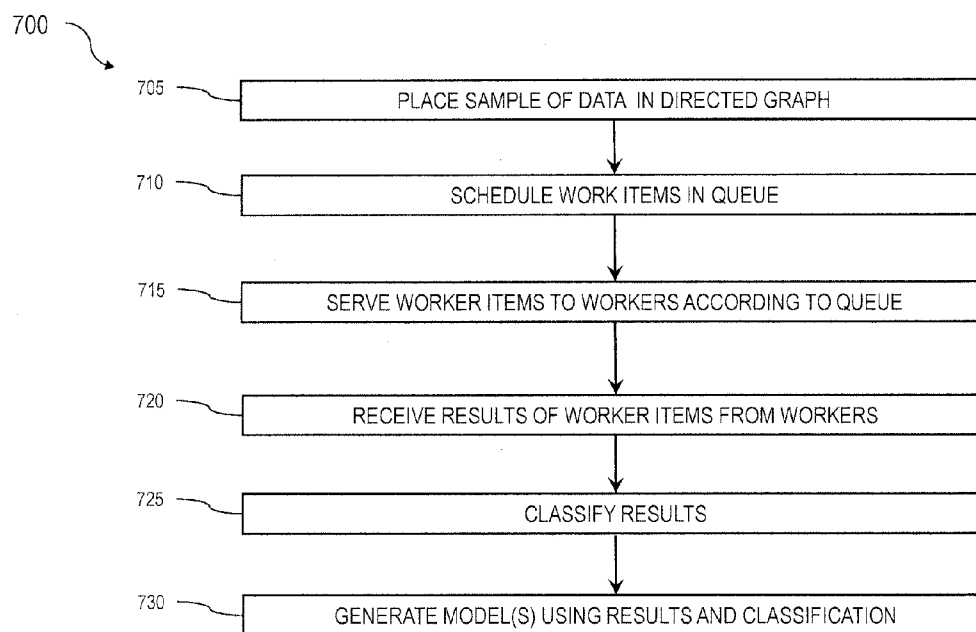


FIG. 7

AUTOMATED SYSTEM FOR GENERATIVE MULTIMODEL MULTICLASS CLASSIFICATION AND SIMILARITY ANALYSIS USING MACHINE LEARNING

RELATED APPLICATION

[0001] The current application claims priority to U.S. Pat. App. Ser. No. 61/838,820 filed on Jun. 24, 2013, the contents of which are hereby fully incorporated by reference.

TECHNICAL FIELD

[0002] The subject matter described herein relates to systems, methods, and computer program products for automated and generative multimodel multiclass classification and similarity analysis using machine learning.

BACKGROUND

[0003] The space of determining if a sample falls into a category and how closely it compares and in which degrees to other samples is a costly and human intensive problem. Traditional methods require humans to make multiple decisions in the process, which adversely affects scalability and repeatability of the process. Additionally, humans are ill equipped to consider data at the scale required to solve difficult problems. Finally, these systems tend to be either overly general and wildly inefficient or overly specific and focused on specific problems.

SUMMARY

[0004] This current subject matter is directed to a process in which computers can be used to efficiently create classifications and similarity analysis using probabilistic models generated through the principles of machine learning. The process does this in an automated way via use of generative models in which samples further train the system and result in iteratively better models that correctly represent the best predictive capabilities expressed within a particular sample population.

[0005] The process can be defined by five major functionality components and the infrastructure required in order to support these functions.

[0006] Query Interface

[0007] Sample collection

[0008] Feature extraction from samples

[0009] Multiclass Sample Classification and Similarity Analysis

[0010] Model Generation

[0011] A sample is any piece of data that you wish to classify or perform similarity analysis against similar samples. A feature is any salient data point that the system measures from a sample. A model is a single or multimodel probability matrix that defines the likelihood of any sample to be classified in a particular class. A multiclass classifier is one that can support classification in two or more classes. A multimodel classifier is one that uses sub models to handle particular intricacies in a complex sample. A generative classifier is one in which samples used to classify may become training material for future analysis.

[0012] In one aspect, a sample of data is placed within a directed graph (e.g., a directed acyclic graph). The directed graph comprises a plurality of hierarchical nodes that form a queue of work items for a particular worker class that is used to process the sample of data. Work items are scheduled

within the queue for each of a plurality of workers by traversing the nodes of the directed graph. Thereafter, the work items are served to the workers according to the queue. Results are then received from the workers for the work items. In this arrangement, the nodes of the directed graph are traversed based on the received results.

[0013] The results can include extracted features from the sample of data. The samples of data and/or extracted features in some cases can be classified. In addition, at least one model (e.g., machine learning model, etc.) can be generated using the extracted features and/or the classification. In other cases, the sample of data can be used to simply classify the results. Data characterizing the classification can be provided in various fashions such as displaying the data, loading the data into memory, storing the data, and transmitting the data to a remote computing system.

[0014] The results can include routing data that is used to determine where to schedule a next subsequent work item in the queue.

[0015] An order of each sample can be prioritized prior to adding such samples to the queue so that each sample is added to the queue according to the prioritized order. The priorities can be based on a pre-defined rate of processing. Prioritization of the at least one sample can be locally adjusted in real-time. The work items can be scheduled in the queue according to at least one of sample prioritization or worker rate.

[0016] The workers to which the work items are served can be part of a pool having a dynamically changing size based on available resources. The available resources can be based on determined supply and demand.

[0017] In one variation, the sample of data includes files for access or execution by a computing system, and wherein the classification indicates whether or not at least one file likely comprises malicious code. In another variation, the sample of data includes medical imaging data and wherein the classification indicated whether or not at least one portion of the medical imaging data indicates a likelihood of an abnormal condition (e.g., cancerous cells, etc.).

[0018] Computer program products are also described that comprise non-transitory computer readable media storing instructions, which when executed one or more data processors of one or more computing systems, causes at least one data processor to perform operations herein. Similarly, computer systems are also described that may include one or more data processors and a memory coupled to the one or more data processors. The memory may temporarily or permanently store instructions that cause at least one processor to perform one or more of the operations described herein. In addition, methods can be implemented by one or more data processors either within a single computing system or distributed among two or more computing systems. Such computing systems can be connected and can exchange data and/or commands or other instructions or the like via one or more connections, including but not limited to a connection over a network (e.g. the Internet, a wireless wide area network, a local area network, a wide area network, a wired network, or the like), via a direct connection between one or more of the multiple computing systems, etc.

[0019] The subject matter described herein provides many technical advantages. For example, the current subject matter provides for the automatic generation of models thereby obviating the need for human generated models and the errors

associated therewith. Furthermore, the current subject matter provides enhanced techniques for classifying data/files for various applications.

[0020] The details of one or more variations of the subject matter described herein are set forth in the accompanying drawings and the description below. Other features and advantages of the subject matter described herein will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

[0021] FIG. 1 is a system diagram illustrating major components of a system for implementing the current subject matter;

[0022] FIG. 2 is a diagram illustrating a process of using a query interface;

[0023] FIG. 3 is a diagram illustrating a process for active collection;

[0024] FIG. 4 is a diagram illustrating a process for a passive collector interface API;

[0025] FIG. 5 is a diagram illustrating a process of extracting features;

[0026] FIG. 6 is a diagram illustrating a process of model generation; and

[0027] FIG. 7 is a diagram illustrating a process for generative multimodel multiclass classification and similarity analysis.

DETAILED DESCRIPTION

[0028] With regard to diagram 100 of FIG. 1, the system described herein can make use of key underlying infrastructure to provide its services, namely to extract features from a sample of data, make classifications regarding some or all of such extracted features, and generate machine learning models using such data that can be used to characterize subsequently data samples.

[0029] The infrastructure can include a unified data access layer 135 that access a cache 140 that stores data from data sources such as relational data stores 145 and big data sources 150. As will be described in further detail below, the infrastructure can include a router 155, a scheduler 160, and a resource manager 165. Unless otherwise specified in this detailed description and/or in the claims, the components of the infrastructure can be implemented in software, hardware, or a combination of both. The infrastructure is shared among the functions and acts as glue that binds the functions together. It has a key focus on optimization of the overarching process. It facilitates data flowing through the system, determines which resources are required, and appropriately schedules an optimal pathing and resource profile for the data.

[0030] The processes described herein can use a dynamically configurable workflow system to define ordering of sample data through an optimal path. An external source 105 (e.g., via a web service and/or other network, etc.) can query data from a query interface 110 and/or submit data to a collection interface 115. Various features from data obtained from an interface can be extracted by an extraction component 120. In some variations, the extracted data can be classified (or otherwise characterized) by a classifier component 125. In addition, a model generation component 130 can generate or otherwise use one or more models (e.g., machine learning models, etc.) based on the extracted and/or classified data.

[0031] With the current subject matter, individual workers within the workflow system do not require specific knowledge of other workers, and do not assume a particular path through the workflow. They can however specify dependencies that will be honored by the scheduler 160 and router 155.

[0032] The workflow can use a centralized but scalable routing and prioritization system. Individual workers register with the central resource manager 165 and receive work via a pull mechanism. In addition, the workflow system can support a high degree of concurrency and flexibility. Work items can be heavy or light weight and will be scheduled to appropriate workers.

[0033] The workflow system can use a system to allow workers of each particular type to be added and removed without restarting the system. It also allows new ordering to be implemented and new types of workers to be added and existing workers to be removed.

[0034] The optimal workflow can be represented as a directed graph such as a directed acyclic graph (DAG), in which nodes are represented as the individual classes of workers. While the current subject matter uses DAGs as an example, other types of hierarchical arrangements/directed graphs can be utilized. The graph is defined by worker class prerequisites in a backward chain dependency mapping process that attempts to generate the shortest path through the graph. This process results in a forward chaining of worker classes that represent the optimal pathing through the workflow.

[0035] Initially, individual worker classes can define prerequisites via a configuration file. This configuration contains a seed of a routing system, allowing an optimized route to be built up at runtime. Thereafter, individual worker classes can be built into a DAG which can then be published to, for example, a central resource manager 165. The central resource manager 165 can then use this new DAG to determine optimal routes for traversing a particular sample of data (which results in selective scheduling of work items to workers/classes of workers).

[0036] The central resource manager 165 using the router 155 can define where work items go next. This frees individual workers from having to know or care how messages are passed. The central resource manager 165 relies on the configured graph paired with the output of each worker to define the optimal next step in the workflow process.

[0037] A sample of data can be placed in a state in the DAG, represented as a queue of work items for a particular worker class. Next, a worker can be served the work item by the router 155. The worker can then return the work item and, in some cases, data useful for routing such as additional layers of information to assist in routing. For instance, if a worker determines that the object has shown signs of corruption, flagging this can help route the object to a better solution. The central manager can now determine the optimal next route based on the results of the current work item and the results of the current worker. The central resource manager 165 then shifts the sample to the next state as defined by the DAG.

[0038] The central resource manager 165 using the scheduler 160 can be responsible for scheduling work for individual workers. The central resource manager 165 can do this by managing the ordering of work items in the queue. For each work item class, represented as a queue of work items, individual items can be provided to workers based on a variety of factors including sample prioritization and worker rate. Additionally, the scheduler 160 is responsible for handling

work items in exceptional cases, such as those being handled by workers who have gone offline or are operating out of parameters.

[0039] In some variations, a sample can enter a particular DAG state. The scheduler **160** can then determine the order of the sample in the queue. The scheduler **160** can watch for work items that have been in a checked out state for too long and reinsert them in the waiting state at the appropriate place.

[0040] Work items can be prioritized by the scheduler **160** based on a rate of processing. Prioritization offers both global and local optimization potential. Local prioritization optimization focuses on operations within a node in the system, where global optimizations involve prioritization operations that affect multiple hosts and route paths. A higher priority sample gets precedence over lower priority samples. Prioritization of a sample can be locally adjusted in real-time to accommodate exceptional scenarios. Higher priority can be reserved for sample processing that require a shorter time to complete analysis. Lower priority can be reserved for bulk processing from passive sources such as backfill operations or samples read from feeds.

[0041] For global prioritization optimization, individual samples can be assigned a prioritization based on the mechanism of insertion into the workflow. The scheduler **160** can use this global prioritization throughout the workflow, giving higher priority work items shorter queue wait times, and lower priority work items longer wait times.

[0042] For local prioritization optimization, the scheduler **160** can determine that the current prioritization for a work item is incorrect or causing a block in the queue. In such cases, the scheduler **160** can dynamically adjust the prioritization for the work item higher or lower depending on the scenario.

[0043] The infrastructure can provide a dynamic scaling factor based on a number of criteria. The primary global metric for scaling can be total processing time for a work item. To adjust a global metric in a heavily distributed system with high variance in sample diversity, a number of local metrics can be considered and adapted to. These local metrics can include, for example: average sample processing time per activity, individual activity work factor, per host load metrics, and available RAM, CPU, and disk. Additionally, the overall system can ensure continuous operation of the whole system in the event of small or large localized failures. In this sense, the scheduler **160** can provide a failsafe self-healing mechanism. The dynamic scaling can be achieved via a combination of detailed monitoring and rapid adaptation of the available resource pool (i.e., the pool of workers, etc.), as well as optionally scaling the complete pool of resources. This process can be managed by the central resource manager **165**.

[0044] The resource management system as implemented by the resource manager **165** can utilize the concept of resource pools. A resource pool is a representation of all computer resources available within a particular class. Classes can generally be defined on commonalities such as operating system, or services offered. Each pool can be divided into resources given to individual workers.

[0045] Resource pools via the resource manager **165** can grow or shrink based on need and resource availability and cost. When resources are readily available and low cost, the resource manager **165** can grow the overall pool and service higher volumes. When resources become scarce or costly, the resource manager **165** can shrink the overall pool and adjust the overall rate to suit requirements.

[0046] The resource management system as implemented by the resource manager **165** can utilize a variety of externally measured and self-reported feedback metrics that help determine resource pool utilization rates, and perceived demand.

[0047] The following are two types of feedback metrics that can be reported to the resource manager **165**. First, measured metrics are those collected by a metric and statistics collection system. They tend to focus on the hard asset metrics such as CPU utilization, available memory or disk space. These are useful to help the resource manager **165** determine resource pool usage metrics, and define under- or over-utilized resources in the pool and adjust the load appropriately. Second, self-reported metrics are useful for determining deeper level of details. The workers and the resource manager **165** report data on processing rates and task times that allow us to determine highly granular resource usage. They are also useful for determining intended demand for resources.

[0048] Individual workers can report a variety of metrics to the resource manager **165** regarding their resource usage and internal timings and counts. These values assist in determination of the source of resource contention. Because metrics are generally collected on both sides of a transaction, the resource manager **165** can determine what is the root cause and what is the symptom, allowing the system to focus our resource optimization on the root cause and alleviate the symptom that way.

[0049] A particularly important set of self-reported metrics comes from the scheduler **160**. The scheduler **160** can track individual worker rates of sample processing and calculate aggregate averages for rates (measured supply). The resource manager **165** also knows the size and wait times for each DAG state work queue (measured demand). These metrics are particularly useful for calculating anticipated total processing time, and determining if there are backups in the workflow. The scheduler **160** can report these metrics to the resource manager **165**.

[0050] When an event, such as a work complete event or a resource limit event, occurs, a corresponding metric is sent by the node generating the event to the resource management system. The resource manager **165** can then create time series based aggregate and individualized measurements based on the self-reported data. These time series data points can be graphed and measured against known baselines. When time series data goes beyond certain points an action can be triggered, such as deploying additional resources into the resource pool.

[0051] In addition, the resource manager **165** can utilize various systems to monitor external states of resources including resource utilization rates as well as availability of resources. These measurements can help determine the optimality of usage. External monitoring can be conducted on a peer basis and also by using dedicated monitoring systems. These systems can continuously poll their intended monitoring targets (such as a resource that cannot self-report its state) and report the results to the resource manager **165**.

[0052] On a timer or an untimed loop, a resource monitor can reach out to a target system. The resource monitor can then report the results of the monitoring action to the resource manager **165**.

[0053] By monitoring the intended supply and demand of resources, the resource manager **165** can determine an optimal deployment of resources. It can ask individual workers to reconfigure themselves to suit the intended demand.

[0054] The resource manager **165** can operate on a timed resource leveling cycle. Once per timed period (say 5 minutes), the amount of used resources is compared to the amount of available resources. Any differences in trends of used and available resources should trigger appropriate remediation such as adding or removing resources from the pool. As reconfiguration of resources does incur overhead, it is imperative that the effect of reconfiguration be managed by the resource manager **165** to ensure the overall system is gaining a net benefit. This cycle can be designed to avoid undue constriction and dilation of resources leading to “flapping”, where resources are continuously being reallocated for whichever task is currently the highest demand.

[0055] The resource manager **165** can also adjust the size of the pool based on not only existing supply and demand, but also external factors like availability and cost.

[0056] The resource manager **165** can use a rules-based approach to appropriately weight each metric in its calculations. The rules can be dynamically configurable but measured and refined based on the existence of available metrics and their value to the overall calculation.

[0057] One process for resource pool reallocation is as follows. Initially, the resource manager **165** can isolate a set of appropriate metrics collected during the cycle period under observation. The set can be defined based on a rules based approach. The resource manager **165** can then determine supply and/or demand from its set of metrics. The resource manager **165** can then examine the current resource pool allocation in light of the existing supply and demand. Next, the resource manager **165** can calculate a new resource pool allocation based on its calculations of supply and demand. In addition, the resource manager **165** can reconfigure workers by asking existing workers to perform more necessary tasks, by adding new workers, or by removing workers altogether, to match the new resource pool allocation. The new allocation pool can then be saved for use by the resource manager **165**/scheduler **160**/router **155** for a next period.

[0058] For resource pool scaling, the resource manager **165** can determine if the current resource pool allocation is too tight or too slack based on a threshold value (e.g., if 80% of workers are busy, >90% of the time, add more workers; or if <10% of workers are busy <50% of the time, remove workers from the pool, etc.). It is too tight if demand is outpacing supply. It is too slack if supply is outpacing demand. If the resource pool is too slack, the resource manager **165** can determine an optimal value to reduce the resource pool. The resource manager **165** can reallocate a new resource allocation for the reduced pool and engage the workers to reconfigure based on this new allocation. If the pool is too tight, the resource manager **165** can examine external cost and availability of additional resources. Next, the resource manager **165** can choose to increase resource pool size, or decrease incoming volume by rate limiting. Rate limiting happens when a global processing time metric is being met. Resource pool increases happen when they are not. Further, the resource manager **165** can reallocate a new resource allocation for the increased pool and engage to reconfigure based on this new allocation.

[0059] The resource manager **165** can also serve the purpose of ensuring operation in adverse conditions. The system can be designed to presume that individual workers are temporal and may appear and disappear both under the control of the resource manager **165** and due to external factors. To ensure continuous operation, the resource manager **165** can

utilize its monitoring system to address problems that arise in the operations. Like resource scaling and optimization, problems can be defined by a specific set of rules applied to metrics. In this case, the metrics are typically external in nature, as an internally reported metric may not be available if the system that reports that metric has a fatal situation.

[0060] The resource manager **165** can decommission continually underperforming assets in the resource pool and replace them with better performing assets. It can also raise alerts for manual intervention in the event that automated responses are not effective in alleviating the problems.

[0061] Self-healing can be implemented by having the resource manager **165** examine the current resource pool for underperforming and non-responsive assets. Thereafter, the resource manager **165** can create a new allocation using existing pool resources or scale the pool larger to replace those resources that are failing. The resource manager **165** can ask the new resources to configure themselves to replace the old resources. The resource manager **165** can later decommission old resources.

[0062] The current system utilizes several types of data stores which can be accessed by the unified data access layer **135**. Various elements of the overall system have different requirements for data consistency, persistence, performance, and integrity. The system can utilize elements of large document based stores (e.g., big data **150**), relational stores **145**, and persistent and temporal caches **140**. To achieve scale and resilience, all of these systems can utilize redundancy and horizontal sharding of data. Additionally, the unified data access layer **135** can have various needs for security levels of the data, as not every element needs complete and unrestricted access to all of the data.

[0063] The infrastructure can also provide a system to unify access to the resources. It can allow each of the systems to provision its particular needs, and it matches these needs with an appropriate backend store. It offers the system the ability to manage access and monitor usage for optimal access. Additionally it can offer an abstraction layer that reduces the requirements needed to implement a worker process.

[0064] The workers and various elements of the system access the unified data access layer **135** through a specified API. The unified data access layer **135** can use a template driven system to allow for rapid mapping of API calls to underlying data sources. The API can be primarily REST based and it can support SSL and multiple forms of authentication.

[0065] Most data access API operations can be managed by a component referred to as a data access manager. The data access manager can support addition and removal of API interfaces, handle management of back-end data resources, and manage security access to these interfaces via the supported mechanisms.

[0066] In order to create a data API, a particular need for a data access API can be defined. Parameters can be defined for consistency, performance, and persistence. Requirements regarding security and confidentiality can also be defined. The API manager can then provision appropriate back end resources to make them available. Thereafter, the API manager can link the back end resources with a front end REST call, with an optional translation operation.

[0067] In order to access an existing data API, the component in question does a service lookup request for the data access API service. Once located, the component can query the API and determine which services it offers. The service

definition can include details on parameterization of access as well as security mechanisms required to access those services and serves as a guide for a component to access an API. When the component knows where a service is located, and how to access it, it can make calls to the data access API as it wishes. In the event of a failover, a redundant peer can serve requests.

[0068] The unified data access layer **135** can support multiple back end solutions. This arrangement helps developers to avoid worrying about access to multiple data sources, using multiple confusing drivers that may or may not be available on the platform. The back-end systems are defined by a set of conditions available to them such as resilience, performance, persistence, consistency, and cost.

[0069] For each backend system, it must be clear at which level it can reliably provide each of the above considerations. Data access can state the desired level needed. Specific levels of acceptable conditions can be stated in the API requirements. This defines which back end systems can effectively serve these requests. In the event that no back end systems can achieve the desired goal, an application can either operate at a degraded level (lower its expectations), or a new back-end data source can be deployed that may meet the criteria in question.

[0070] Resilience can be defined as the ability to continue operations under stress. Any backend component should be able to lose various parts of its operational capacity and be able to adapt effectively to adjust load across the remaining parts. There may be some degree of data loss as a negotiated requirement.

[0071] Performance can be defined by the speed at which data can be delivered or written. Some particularly complex data access operations can be long running and immediate response is not a requirement. Others require specific response times to ensure the viability of the application.

[0072] Persistence can be defined as the amount of time and size of the required data that must be kept in a recoverable state. The period could be from minutes to forever, and the data size requirements are necessary to help properly plan and scale the back-end data stores.

[0073] Consistency can be defined as the assurance that all nodes in a service grouping give the exact same result. This can be affected by certain clustering and replication actions as well as network distance between the clustered nodes. A system with no consistency does not bother with ensuring all nodes are the same. One with high consistency attempts to ensure that all nodes have the same data in near real-time. A common practice is one of “eventual consistency”, in which the system will reach a defined equilibrium of consistency but makes no guarantee on the defined period.

[0074] The query interface **110** is a general component of the overall system that allows external entities to ask questions in regard to samples that have already been processed. The query interface **110** serves as the primary point of contact between external systems **105** and the accumulated knowledge of the system. It is intended to answer questions in a highly efficient manner. The query interface **110** only answers questions for which it has existing answers. Those questions for which the query interface **110** does not have existing answers are passed to the collection interface **115** for further analysis.

[0075] The query interface **110** can utilize a REST API. This allows it to utilize session encryption via SSL and offer a variety of authentication options.

[0076] The query interface **110** can be designed to answer questions based on specific acceptable metadata regarding the sample. This allows questions to be asked without requiring a transfer of the full sample. This can be helpful when the sample is large, complex, or not directly available.

[0077] The query interface **110** can be a specialized layer on top of the unified data access layer **135**. The query interface **110** can interpret external requests via the API into queries against the unified data access layer **135**. For those elements that it cannot answer via this access, the query interface **110** can provide details for the client to access the collection interface **115** to provide a path to get additional information.

[0078] FIG. 2 shows a process flow diagram **200** related to the query interface **110**. Initially, at **205**, an external developer can determine they require the ability to query a specific sample set. The developer can then define the specific metadata required for access. Thereafter, at **210**, the query interface **110** can implement an appropriate set of unified data access layer views into its data stores. The query interface **110** documentation can inform the developer of the specifics of the query interface API endpoint for the sample set they require. It also provides appropriate details in regard to required authentication and encryption. Subsequently, at **215**, the developer can implement the appropriate query interface API within their client.

[0079] As the developer’s program requires access to specific information about a sample, the client, at **220**, can send the appropriate metadata comprising a question. Thereafter, at **225**, the query interface **110** can check the unified data access layer **135** and, at **230**, responds with the answers, or if answers are unavailable, at **235** respond with appropriate details on how to access the collection interface **115** (i.e., ask for a submission, etc.).

[0080] The collection framework components including the collection interface **115** can allow a defined path to get samples of data into the system for analysis and growing our knowledge. These samples serve as bits of information that offer immediate value based on the highest configured model as a classification. They also offer similarity analysis capability, allowing a submitted sample to be put into context among existing samples.

[0081] The collection interface **115** complements the query interface **110** as a primary external conduit for input into the system. The collection interface **115** can allow systems to be interfaced to provide samples for input and receive responses as output. The collection interface **115** can do this via a series of REST API interfaces. These API interfaces can be either public or private, and support SSL and authentication of various sorts.

[0082] Samples can be gathered in one of two general ways. Samples may be “pulled into the system” using an active collection mechanism defined within the system. Alternatively, they may be “pushed” to the system via any collection interface API compatible solution. The combination of active and passive collection allows the system to harvest samples and integrate into existing products in a very light touch way.

[0083] The collection interface **115** can make use of mechanisms to reduce the incidence of unintentional reprocessing. The collection interface **115** can do this by offering basic and complex caching of existing results, and a submission mechanism that can be adjusted to accept or reject submissions based on rules such as current existence and perceived fitness of the sample.

[0084] The collection interface **115** can integrate into the resource manager **165** provided by the general infrastructure. From this integration, the collection interface **115** can adjust rates of input to meet goals, and can offer a prioritization of sample submissions giving precedence on a variety of factors such as preconfigured API priority or dynamic rate adjustment in active collection. This offers a fundamental method to allow for dynamic optimization of submission rates utilizing whole system feedback. It can also help alleviate pressure on existing resource pools in lieu of growing or shrinking existing resource pools. When a sample is collected either passively or actively, the collection interface **115** can push the sample into the appropriate workflow system for processing.

[0085] Active collection offers a process in which the collection framework can manage actively gathering samples for submission. It can offer the ability to manage and schedule collection via the collection interface **115**. It adds the benefit of using feedback from the existing resource manager **165** to dynamically adjust rate if desired.

[0086] An active collector is a subsystem that does some action on a specific period. The active collection framework can manage the rules that define what an active collector may do. Additionally, it can define a mechanism that works with the existing infrastructure scheduling system to support periodic execution of the collectors. This period can be defined as once per timeframe (ex: every 30 minutes), once per clock/calendar time (ex: 1 PM every Monday), or in a continuous manner. The scheduler **160** can ensure appropriate tasks are issued to workers on the correct period.

[0087] The rules that define what a collector can do can be complex and flexible to support a variety of use cases. Collectors can allow internal code to be run on the period, external programs to be run, or even make API calls on the period.

[0088] Active collectors generally do not interface with the query interface **110**; rather they can access the unified data access layer **135** directly. Because these requests do not originate external to the system, they generally have less need for knowing about the current state of a sample and the associated knowledge the system has of that sample.

[0089] FIG. 3 is a diagram **300** illustrating a process for an active collector. Initially, at **305**, a developer can define and build an active collection method. The collection framework can require rules to be written that define the period of active collection (either once per time range, once per set repeated fixed time, or continuous) which the developer, at **310**, can provide. In addition, at **315**, the developer can set rules of collection such as how the active collection is gathered. Once complete, at **320**, the collection framework can manage the scheduler **160** to introduce the task (i.e., apply the active collector) at the appropriate period. Next, at **325**, the scheduler performs the required logic for the active collector on the appropriate period.

[0090] One type of active collection can utilize a web crawling infrastructure. This web crawling infrastructure can allow a general or specialized web crawl to attempt to find suitable samples on the web for analysis. The general web crawler can process HTML pages in a standard manner. It parses the HTML and attempts to locate links and resources that would make appropriate samples, and pushes those samples into the system.

[0091] A specialized web crawl can be one whose scope is limited, and uses specific techniques to make the scan more appropriate. For instance, a specialized web scan can be performed against a specific site, tailored to that site. It can

incorporate site-specific details such as previously gathered information to make the query more effective.

[0092] Another type of active collection can involve running arbitrarily complex queries either in the infrastructure's unified data access layer **135** or in any available external data store. Like all active collection mechanisms, this happens on a specific period.

[0093] When the period defines execution, the active data store collector can reach out to the appropriate data store and collect data to be processed. The active data store can do this individually or via a bulk query.

[0094] Passive collection is a subsystem of the collections interface **115** that can facilitate allowing external access to submitting new samples in an on-demand manner. Generally, it can be used in conjunction to the query interface **110**. Passive collection can provide a specific API that allows samples to be submitted. This API can be generally useful for when existing knowledge of a sample is missing, incomplete, or otherwise unavailable. Such an API can be a REST API that, in turn, enables SSL for session encryption, and offers a variety of authentication mechanisms to ensure desired levels of authentication and access control.

[0095] The passive collection interface API can generally operate in a three-stage manner. First, the passive collection interface API can confirm that the system is ready to accept a submission of a sample. If the system desires the sample, it provides details on where to put the sample. Finally, it expects a confirmation that the sample is put in the proper place.

[0096] The above system allows for a configurable method of upload. The original response from the passive collection interface API defines where the sample should be uploaded, along with any constraints or caveats. This can be a simple HTTP upload, a reference to an FTP or other publicly available URI, or an entirely other system as defined by the sample type and specifics of this instance of the passive collection interface API. This allows the process of uploading to be properly offloaded, and abstracted to best suit multiple types of samples.

[0097] In order to develop a passive collector, developers can consider the elements they wish to receive from an existing product or system. The developer can implement the collection submission client API within their existing product or system. This implementation requires configuration on which submission API endpoint to use, and any credentials if required. During normal product operations, submissions may be generated. Each submission can be sent to the collection submission API. The collection submission API can determine classification and similarity analysis based on current knowledge of the sample and return those results. If additional details are required for analysis, the submission API can offer additional functions to request these details. Failure to provide these details on behalf of the developer may result in lower confidence scores for classification and similarity analysis.

[0098] FIG. 4 is a diagram **400** for accessing the passive collection API. The query interface first provides, at **405**, details on the passive collection API to the client. Thereafter, at **410**, an external client entity asks to submit a sample (i.e., requests a submission to the passive collection interface API. This can be paired internally with the query interface API. Next, at **415**, the passive collection interface system can determine if it desires a submission of the sample based on specifically designated criteria (previous existence, age of existing knowledge, etc.). If the passive collection interface

system desires the sample, it responds to the external client entity with an affirmative response and details on where to “put” the sample, otherwise, at **420**, the client receives a negative response. Depending on the sample type, this may be one of a variety of options. The external client entity can perform, at **425**, the necessary actions to “put” the sample in the correct place. The external client entity can then submit, at **430**, a confirmation to the passive collection interface API. Upon a successful confirmation, the passive collection interface system can, at **435**, place the sample in the proper workflow for further processing.

[0099] The extraction component **120** can serve the purpose of converting samples into salient data points, or “features” used for classification and similarity analysis. The extraction component **120** can do this by using a highly scalable, parallel processing distributed system to extract this information. The extraction component can be defined by a series of worker classes handled through the infrastructure dynamic workflow system. These worker classes can be developed on a per-sample type basis and allow the system to grow in a scalable manner. These worker classes can be effectively chained together by the DAG router **155** in the workflow system to ensure that additional extraction workers can be added and removed at runtime in a dynamic manner.

[0100] Each worker class can represent some set of extracted features and can be associated with a specific sample type. The system can use dynamic typing of samples to ensure that the correct operations are performed on a per sample basis. It achieves this via the use of the infrastructure router **155** and scheduler **160**. The router **155** can ensure that samples are processed in an appropriate manner. The scheduler **160** can ensure that the samples are processed in a timely manner.

[0101] The resource manager **165** can ensure that there are adequate resources available to perform each class of work in an on-demand and planned basis. Each worker class represents a resource that needs to be managed by the resource manager **165**.

[0102] The worker class registers with the router **155** and other infrastructure systems to help the system understand the capabilities it has on hand. When a worker class registers, it provides a set of features it can provide to the router **155**, as well as a set of routes and associated prerequisites and conditionals for each route. Each route can either be a default route (one taken if no other route has been selected), or a named route. Each named route has a series of conditions that must be met to be selected. If these conditions are met (based on the features accumulated for a sample), this named route is selected. There can be multiple named routes that have differing conditions, allowing for a complex routing scenario. The precedence of routing is inherent in the ordering of the routes upon creation.

[0103] The extraction component **120** can use a dynamic set of routing conditions that are dynamically defined in process. As a sample navigates the workflow as assigned by the router **155**, the extraction component **120** can add additional information to the sample profile, allowing for more specialized worker class operations to be performed that produce deeper details.

[0104] FIG. 5 is a diagram **500** illustrating extraction of features. Initially, at **505**, the router **155** can place a work item in a worker queue. Each work item can comprise one or more tasks to effect extraction. Subsequently, at **510**, the worker receives a work item from the worker queue and, at **515**,

processes it. If the task was successfully processed, it is determined, at **520**, whether there are steps/tasks to perform. If the task was not successfully processed, the work item (or a portion thereof) can be placed back in the worker queue and the process repeats. If it is determined that there are no new steps (at **520**), then, at **525**, extraction is finished, otherwise the router **155** sends the work item to the next step. A task is a bit of work performed on a sample, a step is a route to the next task to be performed. In graph theory, a task would be considered a node (or vertex), while a route would be an edge.

[0105] This process can occur with a dynamic typing system that adjusts based on previous knowledge collected in process. The existence of this knowledge is one of the factors used to create the reverse chain Directed Acyclic Graph used by the router. This ensures that the information required to make a decision at any point in the workflow is available when needed.

[0106] This arrangement generally defines a set of prerequisites and conditional factors on a per work class basis. The prerequisites define those data points that need to exist before the routing decision can be made by the router **155**, and the conditionals are those features that are salient to the routing decision. Additionally, conditionals and can be performed on non-feature based data, such as an external API call.

[0107] A dynamic type-based route can be specified by a developer first defining those conditions that are required to perform the appropriate decision. Thereafter, the developer can determine which features can provide the data needed to meet the conditionals and mark these as prerequisites. The developer later can define the features this work class will provide. The router logic system can then calculate a Directed Acyclic Graph based on the required prerequisites, conditionals, and provided data points. The router logic system can build optimal models for the workflow based on this data.

[0108] Worker classes can form emergent “chains” of operations based on the Directed Acyclic Graph built and maintained by the router **155**. The chain can define the operations required to optimally extract all available features from a particular sample. While the workflow is defined in process, some chains consistently match certain types of samples.

[0109] A chain can be locally optimized to reduce the number and types of transitions. The resource manager **165** can do this by allocating partial and whole chains to a single compute resource, allowing operations to happen in an effective manner with minimal transfer time. This arrangement can reduce the overhead involved in forcing data to be moved around in a parallel-distributed system, greatly enhancing the throughput the system provides. This arrangement can also allow for the creation of specific enclaves of processing, which is useful if a particular sample type has an anomalous usage structure (i.e. it uses a significant amount of resources or takes a long time to process).

[0110] Chains can be continually measured and optimized. This arrangement can allow for both automatic and manual adjustments based on ongoing feedback from the system. Knowledge of how chains perform offers the resource manager **165** better knowledge to optimize and anticipate resource utilization based on probability of anticipated chain usage. The resource manager **165** can also utilize real-time chain metrics to define if specific chains are over- or under-performing and adjust its global resource usage allocation appropriately. Chains that are infrequently used can be reduced and have resources assigned to those chains that are in heavy use.

[0111] New chains can be continually added to both consider new sample types as well as add additional depth to existing sample types. With this, the overall system gains new sources of information available to refine and consider new models for classification and similarity analysis of samples.

[0112] The classifier component 125 can provide a method to classify features extracted from a sample against existing models. The classifier component 125 can also provide similarity analysis across defined similarity points. The classifier component 125 can operate as independent worker classes focused on creating classifications. Like the extraction component 120, the classifier component 125 can create emergent chains within a Directed Acyclic Graph that the router 155 can use to push samples through via routing decisions. In addition, like extractors, these chains can be specific to sample types, allowing for a large amount of specialization in classification. Some of this specialization will be in regard to specific feature sets associated with individual samples.

[0113] The classification chain can comprise a set of workers that operate similarly to the extraction workers, or analysis workers. Their purpose can be to specifically create classifications and perform additional classification logic based on internal data like features and external data like data from exterior classification systems. Generally, the internal classifications can be based in part or whole on the output of specialized classifiers embodying the machine learning algorithms. Even those that do not utilize machine learning directly benefit from it, as it is used to measure optimal aggregation algorithms to best represent the current state of knowledge of the model.

[0114] The classifier component 125 can utilize the router 155 and scheduler 160 and abide by all of the inherent aspects of prioritization and optimization of routes. Additionally, the classifier component 125 can operate within the resource manager 165, allowing rates to be adjusted in real-time by allocating more resources to the classifier worker classes.

[0115] Like classification, similarity analysis can be performed using analytics models trained against a representative sampling of the population of all samples. Similar chains can handle large scale directed and undirected similarity analysis. Generally, directed similarity analysis can offer a higher degree of useful output. Undirected similarity analysis tend to be used more to explore sameness against the population, and to attempt to identify those samples that are fundamentally different from anything previously seen. Each similarity analysis can be time and model dependent. Having a new model or a different set of samples in the population may cause the similarity analysis to be recalculated.

[0116] Similarity analysis can further allow the detection of suitably anomalous “outliers”. These outliers can be useful to determine that the models in question are still actively in range. If the rate of outlier detection goes up, the models being used can be re-evaluated to incorporate the newest types of samples within the population.

[0117] Classification can be conducted by the router 155 deciding that a sample is ready for classification when all requisite elements have been collected or loaded from a previous data store. Thereafter, the router 155 can shift the sample into the head of an appropriate classifier chain. The chain can then be executed and a final classification score can be calculated. The classifier chain can interface with the unified data access layer 135 so that the results of the classification can be stored. The classification of the sample is now available for future queries and usage in other places.

[0118] Similarity analysis can be conducted by the router 155 deciding that a sample is ready for similarity analysis when all requisite elements have been collected or loaded from a previous data store. The router 155 then shifts the sample into the head of an appropriate similarity analysis chain. The chain can then be executed and sets of similarity analysis scores can be calculated. Undirected similarity analysis can be stored in an area reviewable by humans to infer meaning. The similarity analysis chain can interface with the unified data access layer 135 so that the results of the directed similarity analyses can be stored. The similarity of the sample can then be available for future queries and usage in other places.

[0119] The classifier work items can support the use of aggregate classifications. These types of classifications can be layers built on top of previously calculated classifications. In a first pass classification, a sample can receive several individualized classification scores. A second stage can allow additional logic to be defined to best calculate effective aggregate classification scores to suit a particular need.

[0120] The use of a multi-stage classifier can allow complete and partial reclassification at high velocity and volume with minimal resource utilization. By leveraging the scheduler and router, missing elements in the multi-stage classifier can be backfilled if they are missing, incomplete, or not up to date. These calculations can be performed during processing, offering the use of the latest and most complete picture on each classification decision.

[0121] The system can use models to answer classification and similarity analysis questions about samples. Such models that drive the classifiers can be defined, developed and built. For these models to be effective, the cost of model generation needs to be minimized and optimized for cost and accuracy. A model can comprise a probability matrix based on machine learning techniques. To do this, a training set of samples can be extracted from a sample population and used to generate a model (i.e., a model can be trained using historical data in order to characterize future data sets/samples, etc.). The model can then be back-tested against a large validation set (that does not contain the sample set). Once a model is judged valid, it can be put into production use for the classifiers.

[0122] A subset of the total population can, in most cases, be used because the samples have a reasonably standard distribution. Because of this, using a large enough total population, and a large enough subset of that for training picked at random, results in a fair representation of the samples in total. Due to ongoing submissions of samples into the system, the number of sample sets and training sets can continuously increase. This arrangement allows iterative models to be built that have enough density to allow sub models to be built to perform better analysis in specialized classifiers.

[0123] While it is helpful to consider all samples as functionally equivalent, the reality is that there are logical gradients when moving from very general similarities to very specific similarities. The system can examine and attempt to create useful sub models to further refine these segregations. The models themselves can result in a multi-stage classifier that can allow a base level of very general features to be compared, and can result in more specific features to be compared against generally similar types of samples.

[0124] The model generation system can be designed to be hands off, allowing the use of internal measures to attempt to create optimal models. This involves defining success and

failure criteria and finding the way to measure and compare the fitness of models to a particular task.

[0125] FIG. 6 is a diagram 600 that illustrates a process for general model generation. Initially, at 605, a complete census of the sample population can be taken (which can be referred to as the “whole population”). Thereafter, at 610, a random subset can be selected to act as a training set. In addition, at 615, a second random subset of the population can be selected from the remaining whole to serve as a testing set. Next, at 620, a model can be generated by applying machine learning classification techniques to the training set. The training set is processed into large vectors of numbers, and these vectors are used in a variety of machine learning algorithms including logistical regression, neural networks, support vector machines, and decision tree ensembles with custom tuned variables to generate a series of models. The models are further tested and refined to produce end stage models for use in the general system. Thereafter, at 625, the generated model can be tested against the testing set and, optionally, confidence intervals can be calculated. The confidence intervals can be compared against existing models to judge fitness and make a determination whether the model failed which causes 610-625 to be repeated or whether to publish the model (at 630).

[0126] The training set generator can be configured to attempt to reduce any specific bias in the selection of samples for the training set from the general sample population. The training set generator can do this by randomly selecting training samples from the population.

[0127] Sub models presuppose a type of bias when compared against the general set, and so it is important that sub models only be compared against those appropriate samples for which they can be representative. This is because a sub model is created against a subset of the entire population. To generate sub models, it is important to apply certain criteria as a primary filter. As long as this primary filter is applied uniformly, both for sample selection and for future classification, and samples are chosen at random from the post filter results, there is no additional bias introduced, at least within the scope of items pertaining to the sub model.

[0128] As the distribution of samples may not be exactly “normal”, from a statistical purpose, an individual training set may not be entirely representative. To combat this, the system can attempt to increase the training set size and run multiple iterations to ensure that the variance of a particular model is not unduly biased.

[0129] The iterative models can add additional weightings to favored features based on appropriateness. By increasing the number of iterations, confidence of the validity of the weights of a feature set can be gained, resulting in a stronger overall classifier.

[0130] Models can use some subset of available features to generate classification models. Because the number of salient features can become rather large over time, and the training sets may become rather large, there may be a desire to reduce the set of features considered in an attempt to optimize the generation process. The process of pruning features can be done in an iterative manner. By reducing the set of features to those that positively affect the model, it can drastically reduce overhead in classification. Each iteration, those features that provide weak scoring (as based on a fitness function) can be dropped in subsequent iterations. This is an optional step, as full models may be more accurate than pruned/reduced models.

[0131] The set of observed features for a sample can be a subset of the overall feature space, due to either the absence of a subset of features from a sample, or the inability to observe a complete set of features. To allow for statistical comparisons between samples of various feature subsets, the system can, for instance, generate estimated features for those features not in a sample subset, define absent features as valid features for comparison, or utilize a multimodel approach to compare samples at an abstract level, where feature spaces are not directly compared.

[0132] As noted above, models can be effectively trained to use any form of supervised learning. To do this, a set of data for which you know or suspect a classification at a high confidence must be compiled. This training can be done by various methodologies. It can be hand validated, or machine generated in some contexts.

[0133] Unsupervised machine learning does not require training, but can be much more variable in its output. If the system is designed to be exploratory in nature, the training set can be judged to be the entire set.

[0134] Similarity analysis does not necessarily require training. It can take advantage of a previously classified sample to help categorize the similarities in question and judge fitness of those samples, but even this is not required. Classification offers the ability to layer context on top of similarities that may be detected, versus having to determine context from observed similarities.

[0135] Models can be validated based on a level of fitness. Fitness in this sense is the accuracy of the model, along with other facts. The system can allow for the definition of a set of criteria that defines fitness. It can be used as a low-water mark to ensure that the models produced by the system meet minimum criteria. Additional fitness criteria can be garnered by comparisons of the same data against previous models. Generally, fitness can be defined as the accuracy of classification. This accuracy can be considered by the measure and type of correct and incorrect answers. This is a more complex operation when multiple classes are considered, as the system considers these rates in a relative manner, and may choose to perform degrees of closeness to intended results in addition to the standard accuracy tests

[0136] A sample can be classified in one of four categories (calculated per sample on a per class basis):

[0137] 1. True positive—the sample belongs in the classification, and the model placed it there;

[0138] 2. True Negative—the sample does not belong in the classification, and the model did not place it there;

[0139] 3. False Positive—the sample does not belong in the classification, and the model placed it there;

[0140] 4. False Negative—the sample belongs in the classification, and the model did not place it there.

[0141] The goal is to maximize true positives on a per class basis. The rate of error considers the tolerance for false positives and negatives. Various adjustments to the model can be used to improve or further optimize the reduction of these levels of error.

[0142] Once a model has been generated that meets the criteria for fitness, an additional back test can be calculated across the entire population. This back-test can provide a more comprehensive analysis of the samples versus the entire population. If the complete back-test demonstrates a similar fitness characteristic to the model as calculated in the training set, and the fitness meets criteria, at this point, the model can be confirmed. Once a model is confirmed, the classification

for each sample in the population can be updated via the unified data access layer 135 and the new model can be put into active participation via the classifier component 125.

[0143] A general model requires calculation of several smaller specific case models. These models can be defined as part of the overall model during its original genesis. The sub models operate on a very similar overall process to the main model, only on smaller data sets and sets of features. The general model can accommodate multiple sub models for particular specialized criteria. Each sub model should be generated in a similar iterative manner. The final results of the sub models should be referenced in the main model, unless a sub model is judged to be entirely complete in its assessment, in which case it supplants the main model for classifications of this type.

[0144] Multiple models can be generated as follows. For each sub model required, an appropriate entire population can be selected. Thereafter, a random training set can be selected from the sub model population. In addition, a random testing set can be selected from the sub model population. Next, a model can be generated based on the appropriate training set. The fitness of the new sub model can be iteratively tested against defined criteria and previous models. Upon achieving an accurate model, the model can be published in conjunction with the general model.

[0145] FIG. 7 is a diagram 700 of a process flow diagram in which, at 705, a sample of data is placed within a directed graph such as directed acyclic graph. The directed graph comprises a plurality of hierarchical nodes that form a queue of work items for each particular worker class that is used to process the sample of data. Subsequently, at 710, work items are scheduled within the queue for each of a plurality of workers by traversing the nodes of the directed graph. The work items are then served, at 715, to the workers according to the queue. Results can later be received, at 720, from the workers for the work items (the nodes of the directed graph are traversed based on the received results). In addition, in some variations, at 725, the results can be classified so that, at 730, one or models can be generated.

[0146] It will be appreciated that the current subject matter can be utilized across many different applications in which there is a need to classify or otherwise characterize data. In one example, this system can be used to make a determination of the likelihood that a particular computer file is malicious (intends harm to the operator or underlying computer system). In this scenario, the system can be defined by samples representing files on a computer. These files can be normal program executables, data files, or any other type of file on a computer. The classification system is tuned (using, for example, one or models that have been trained using historical file analyses with known outcomes) to model the “goodness” and “badness” of a potential sample, delivering a likelihood of a sample being something that could cause harm to a computer if executed. The models are created from a set of features extracted from samples (using, for example, one or models that have been trained using historical file analyses with known outcomes). These features can include measurement about the file as well as its contents through several stages of analysis. Some example features include file size, information density, structured layout of the file, specific elements pertinent to the type of file it is (program section names for programs, author details for documents, etc.). The features can also include several layers of deeper analysis that can be represented as features within the system. This can

include deep textual or code analysis, including interpreting instructions in an emulated fashion.

[0147] Another example of the use of the current subject matter is to solve image classifications within a biomedical application, such as X-Ray processing. Sample classifications could be the likelihood of the presence of a cancerous growth within an image. In this configuration of the system, samples are represented by individual X-Ray images contained in high resolution image formatted computer files. These images are processed, collecting feature based data including orientation, size, shade differences, and linearity. These features are used to create models that offer strong ability to predict the existence or non existence of cancerous growths being detected in a particular image, and highlight those growths to a researcher or doctor.

[0148] One or more aspects or features of the subject matter described herein can be realized in digital electronic circuitry, integrated circuitry, specially designed application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs) computer hardware, firmware, software, and/or combinations thereof. These various aspects or features can include implementation in one or more computer programs that are executable and/or interpretable on a programmable system including at least one programmable processor, which can be special or general purpose, coupled to receive data and instructions from, and to transmit data and instructions to, a storage system, at least one input device, and at least one output device. The programmable system or computing system may include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[0149] These computer programs, which can also be referred to programs, software, software applications, applications, components, or code, include machine instructions for a programmable processor, and can be implemented in a high-level procedural language, an object-oriented programming language, a functional programming language, a logical programming language, and/or in assembly/machine language. As used herein, the term “machine-readable medium” refers to any computer program product, apparatus and/or device, such as for example magnetic discs, optical disks, memory, and Programmable Logic Devices (PLDs), used to provide machine instructions and/or data to a programmable processor, including a machine-readable medium that receives machine instructions as a machine-readable signal. The term “machine-readable signal” refers to any signal used to provide machine instructions and/or data to a programmable processor. The machine-readable medium can store such machine instructions non-transitorily, such as for example as would a non-transient solid-state memory or a magnetic hard drive or any equivalent storage medium. The machine-readable medium can alternatively or additionally store such machine instructions in a transient manner, such as for example as would a processor cache or other random access memory associated with one or more physical processor cores.

[0150] To provide for interaction with a user, one or more aspects or features of the subject matter described herein can be implemented on a computer having a display device, such as for example a cathode ray tube (CRT) or a liquid crystal display (LCD) or a light emitting diode (LED) monitor for

displaying information to the user and a keyboard and a pointing device, such as for example a mouse or a trackball, by which the user may provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well. For example, feedback provided to the user can be any form of sensory feedback, such as for example visual feedback, auditory feedback, or tactile feedback; and input from the user may be received in any form, including, but not limited to, acoustic, speech, or tactile input. Other possible input devices include, but are not limited to, touch screens or other touch-sensitive devices such as single or multi-point resistive or capacitive trackpads, voice recognition hardware and software, optical scanners, optical pointers, digital image capture devices and associated interpretation software, and the like.

[0151] In the descriptions above and in the claims, phrases such as “at least one of” or “one or more of” may occur followed by a conjunctive list of elements or features. The term “and/or” may also occur in a list of two or more elements or features. Unless otherwise implicitly or explicitly contradicted by the context in which it used, such a phrase is intended to mean any of the listed elements or features individually or any of the recited elements or features in combination with any of the other recited elements or features. For example, the phrases “at least one of A and B;” “one or more of A and B;” and “A and/or B” are each intended to mean “A alone, B alone, or A and B together.” A similar interpretation is also intended for lists including three or more items. For example, the phrases “at least one of A, B, and C;” “one or more of A, B, and C;” and “A, B, and/or C” are each intended to mean “A alone, B alone, C alone, A and B together, A and C together, B and C together, or A and B and C together.” In addition, use of the term “based on,” above and in the claims is intended to mean, “based at least in part on,” such that an unrecited feature or element is also permissible.

[0152] The subject matter described herein can be embodied in systems, apparatus, methods, and/or articles depending on the desired configuration. The implementations set forth in the foregoing description do not represent all implementations consistent with the subject matter described herein. Instead, they are merely some examples consistent with aspects related to the described subject matter. Although a few variations have been described in detail above, other modifications or additions are possible. In particular, further features and/or variations can be provided in addition to those set forth herein. For example, the implementations described above can be directed to various combinations and subcombinations of the disclosed features and/or combinations and subcombinations of several further features disclosed above. In addition, the logic flows depicted in the accompanying figures and/or described herein do not necessarily require the particular order shown, or sequential order, to achieve desirable results. Other implementations may be within the scope of the following claims.

What is claimed is:

1. A method for implementation by one or more data processors forming part of at least one computing system, the method comprising:

placing a sample of data within a directed graph that comprises a plurality of hierarchical nodes that form a queue of work items for a particular worker class that is used to process the sample of data;

scheduling work items within the queue for each of a plurality of workers by traversing the nodes of the directed acyclic graph;
serving the work items to the workers according to the queue; and
receiving results from the workers for the work items;
wherein the nodes of the directed graph are traversed based on the received results.

2. The method of claim 1 wherein the results comprise extracted features from the sample of data.

3. The method of claim 2 further comprising:

classifying the sample of data and/or the extracted features.

4. The method of claim 3 further comprising:

generating at least one model using the extracted features and the classification.

5. The method of claim 1 further comprising:

classifying the sample of data based on the received results.

6. The method of claim 5 further comprising:

providing data characterizing the classifying, the providing comprising at least one of: displaying the data characterizing the classifying, storing the data characterizing the classifying, loading the data characterizing the classifying into memory, or transmitting the data characterizing the classifying to a remote computing system.

7. The method of claim 1, wherein the results further comprise routing data that is used to determine where to schedule a next subsequent work item in the queue.

8. The method of claim 1 further comprising:

prioritizing an order of each sample prior to adding such samples to the queue, wherein each sample is added to the queue according to the prioritized order.

9. The method of claim 8, wherein the priorities are based on a pre-defined rate of processing.

10. The method of claim 8, wherein prioritization of the at least one sample is locally adjusted in real-time.

11. The method of claim 1, wherein the work items are scheduled in the queue according to at least one of sample prioritization or worker rate.

12. The method of claim 1, wherein the workers to which the work items are served are part of a pool having a dynamically changing size based on available resources.

13. The method of claim 12, wherein the available resources are based on determined supply and demand.

14. The method of claim 5, wherein the sample of data comprises files for access or execution by a computing system, and wherein the classification indicates whether or not at least one file likely comprises malicious code.

15. The method of claim 5, wherein the sample of data comprises medical imaging data and wherein the classification indicated whether or not at least one portion of the medical imaging data indicates a likelihood of an abnormal condition.

16. The method of claim 1, wherein the directed graph is a directed acyclic graph.

17. A non-transitory computer program product storing instructions which, when executed by at least one data processor forming part of at least one computing system, result in operations comprising:

placing a sample of data within a directed graph that comprises a plurality of hierarchical nodes that form a queue of work items for a particular worker class that is used to process the sample of data;

scheduling work items within the queue for each of a plurality of workers by traversing the nodes of the directed acyclic graph;
 serving the work items to the workers according to the queue;
 receiving results from the workers for the work items comprising extracted features; and
 classifying at least a portion of the extracted features using one or more machine learning models;
 wherein the nodes of the directed acyclic graph are traversed based on the received results.

18. The computer program product of claim **17**, wherein the sample of data comprises files for access or execution by a computing system, and wherein the classification indicates whether or not at least one file likely comprises malicious code.

19. The computer program product of claim **17**, wherein the sample of data comprises medical imaging data and wherein the classification indicated whether or not at least one portion of the medical imaging data indicates a likelihood of an abnormal condition.

20. A system comprising:
 at least one data processor; and
 memory storing instructions which, when executed by the at least one data processor, result in operations comprising:
 placing a sample of data within a directed acyclic graph, the directed acyclic graph comprising a plurality of

hierarchical nodes that form a queue of work items for a particular worker class that is used to process the sample of data;

scheduling work items within the queue for each of a plurality of workers by traversing the nodes of the directed acyclic graph;

serving the work items to the workers according to the queue;

receiving results from the workers for the work items comprising extracted features;

classifying at least a portion of the extracted features; and

generating at least one machine learning model using the classified extracted features.

wherein the nodes of the directed acyclic graph are traversed based on the received results.

21. The system of claim **20**, wherein the sample of data comprises files for access or execution by a computing system, and wherein the classification indicates whether or not at least one file likely comprises malicious code.

22. The system of claim **20**, wherein the sample of data comprises medical imaging data and wherein the classification indicated whether or not at least one portion of the medical imaging data indicates a likelihood of an abnormal condition.

* * * * *