



- (51) International Patent Classification:
G06F 21/60 (2013.01) G06N 3/02 (2006.01)
- (21) International Application Number:
PCT/US2014/050604
- (22) International Filing Date:
11 August 2014 (11.08.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
61/864,274 9 August 2013 (09.08.2013) US
- (71) Applicant: BEHAVIORAL RECOGNITION SYSTEMS, INC. [US/US]; 2100 West Loop South, 9th Floor, Houston, Texas 77027 (US).
- (72) Inventors: COBB, Wesley Kenneth; 215 Wisteria Walk Circle, The Woodlands, Texas 77381 (US). SEOW, Ming-Jung; 4606 Tilbury Trail, Richmond, Texas 77407 (US). COLE, Curtis Edward, Jr.; 5523 Rivertree Lane, Spring, Texas 77379 (US). FALCON, Cody Shay; 26411 Red Cliff Ridge, Katy, Texas 77494 (US). KONOSKY, Benjamin A.; 142 Moss Point, Friendswood, Texas 77546 (US). MORGAN, Charles Richard; 6002 S. Trafalgar Ct., Katy, Texas 77449 (US). POFFENBERGER, Aaron; 2100 West Loop South, 9th Floor, Houston, Texas 77027 (US). NGUYEN, Thong Toan; 20715 Sommerset Branch Ct., Richmond, Texas 77407 (US).

(74) Agents: PATTERSON, B., Todd et al.; Patterson & Sheridan, L.L.P., 24 Greenway Plaza, Suite 1600, Houston, Texas 77046 (US).

(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, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, 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, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, 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))

(54) Title: COGNITIVE INFORMATION SECURITY USING A BEHAVIOR RECOGNITION SYSTEM

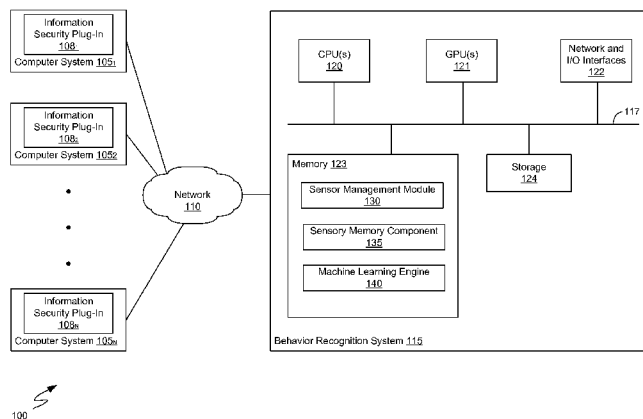


Figure 1

(57) Abstract: Embodiments presented herein describe a method for processing streams of data of one or more networked computer systems. According to one embodiment of the present disclosure, an ordered stream of normalized vectors corresponding to information security data obtained from one or more sensors monitoring a computer network is received. A neuro-linguistic model of the information security data is generated by clustering the ordered stream of vectors and assigning a letter to each cluster, outputting an ordered sequence of letters based on a mapping of the ordered stream of normalized vectors to the clusters, building a dictionary of words from the ordered output of letters, outputting an ordered stream of words based on the ordered output of letters, and generating a plurality of phrases based on the ordered output of words.

WO 2015/021484 A1

COGNITIVE INFORMATION SECURITY USING A BEHAVIOR RECOGNITION SYSTEM

CROSS-REFERENCE TO RELATED APPLICATION

5 [0001] This application claims benefit of U.S. Provisional Patent Application Serial No. 61/864,274, filed August 9, 2013, which is incorporated herein by reference in its entirety.

BACKGROUND

Field

10 [0002] Embodiments presented herein provide techniques for analyzing streams of computer data. More specifically, techniques are disclosed for analyzing and learning patterns of behavior from monitoring computer networks.

Description of the Related Art

15 [0002] Intrusion detection systems (IDS) monitor network or system activities for malicious activities or violations and produce reports to a management console. Many current intrusion detection systems are signature-based. That is, an IDS may be configured with signatures to detect malicious or unwanted activity. As known, an attack signature is a sequence of computer activities (or alterations of those activities) corresponding to a known attack, such as towards a vulnerability in an operating
20 system or application. An attack signature may also be a bit pattern in a network packet corresponding to a known vulnerability. An IDS may use attack signatures through passive protocol analysis (e.g., using a "sniffer" in promiscuous mode) or signature analysis (e.g., interpreting a specific series of packets or data within that packet to a known pattern of attack). For example, an IDS may be configured with an
25 attack signature that detects a particular virus in an email. The signature may contain information about subject field text associated with previous e-mails that have contained the virus or attachment filenames in the past. With the signature, the IDS can compare the subject of each e-mail with subjects contained in the signature and also attachments with known suspicious filenames.

30 [0003] However, the signature-based approach raises several concerns. For instance, although an IDS may possibly detect alterations to a particular attack, the

alterations typically must be defined in the signature to do so. Similarly, because attack signatures are predefined, the IDS is susceptible to new attacks that have not yet been observed. That is, someone generally has to observe a particular attack pattern or instance of an exploit before defining a signature. As a result, an IDS may be unable to detect so-called “zero-day” attacks (i.e., attacks that exploit a previously unknown vulnerability in a system or application). As the number of systems is increasingly being breached by zero-day attacks, the signature-based approach in intrusion detection systems becomes of limited usefulness. In addition, attackers may use code morphing techniques to subvert attack signatures, making the attacks difficult to detect.

SUMMARY

[0004] One embodiment presented herein includes a method for processing streams of data of one or more networked computer systems. The method generally includes receiving an ordered stream of normalized vectors corresponding to information security data obtained from one or more sensors monitoring a computer network. A neuro-linguistic model of the information security data is generated by clustering the ordered stream of vectors and assigning a letter to each cluster, outputting an ordered sequence of letters based on a mapping of the ordered stream of normalized vectors to the clusters, building a dictionary of words from of the ordered output of letters, outputting an ordered stream of words based on the ordered output of letters, and generating a plurality of phrases based on the ordered output of words.

[0005] Another embodiment includes a computer-readable storage medium having instructions, which, when executed on a processor, performs an operation for processing streams of data of one or more networked computer systems. The operation generally includes receiving an ordered stream of normalized vectors corresponding to information security data obtained from one or more sensors monitoring a computer network. A neuro-linguistic model of the information security data is generated by clustering the ordered stream of vectors and assigning a letter to each cluster, outputting an ordered sequence of letters based on a mapping of the ordered stream of normalized vectors to the clusters, building a dictionary of words from of the ordered output of letters, outputting an ordered stream of words based on

the ordered output of letters, and generating a plurality of phrases based on the ordered output of words.

5 [0006] Still another embodiment includes a system having a processor and a memory storing a program, which, when executed on the processor, performs an operation for processing streams of data of one or more networked computer systems. The operation generally includes receiving an ordered stream of normalized vectors corresponding to information security data obtained from one or more sensors monitoring a computer network. A neuro-linguistic model of the information security data is generated by clustering the ordered stream of vectors and assigning a letter to
10 each cluster, outputting an ordered sequence of letters based on a mapping of the ordered stream of normalized vectors to the clusters, building a dictionary of words from of the ordered output of letters, outputting an ordered stream of words based on the ordered output of letters, and generating a plurality of phrases based on the ordered output of words.

15 [0007] Yet another embodiment includes a method for processing streams of data of one or more networked computer systems. The method generally includes receiving a stream of normalized vectors corresponding to information security data monitored over a computer network. A plurality of symbols from the stream of normalized vectors is generated. A dictionary of words from combinations of the
20 generated plurality of symbols is generated based on an order by which symbols appear in the stream relative to one another. A plurality of phrases is generated based on relationships identified between words in the dictionary. Each relationship is based on an order by which each word appears relative to one another. Patterns of behavior are identified within the information security data.

25 **BRIEF DESCRIPTION OF THE DRAWINGS**

[0008] So that the manner in which the above recited features of the present disclosure can be understood in detail, a more particular description of the disclosure, briefly summarized above, may be had by reference to embodiments, some of which are illustrated in the appended drawings. It is to be noted, however, that the
30 appended drawings illustrate only exemplary embodiments and are therefore not to be considered limiting of its scope, may admit to other equally effective embodiments.

[0009] Figure 1 illustrates an example computing environment, according to one embodiment.

[0010] Figure 2 further illustrates the information security plug-in described in Figure 1, according to one embodiment.

5 [0011] Figure 3 illustrates a system architecture of the behavior recognition system described in Figure 1, according to one embodiment.

[0012] Figure 4 illustrates a method for normalizing data sent by an information security plug-in to an information security driver in the behavior recognition system, according to one embodiment.

10 [0013] Figure 5 illustrates a method for generating a linguistic representation of normalized information security data, according to one embodiment.

[0014] To facilitate understanding, identical reference numerals have been used, where possible, to designate identical elements that are common to the figures. It is contemplated that elements and features of one embodiment may be beneficially
15 incorporated in other embodiments without further recitation.

DETAILED DESCRIPTION

[0015] Embodiments presented herein provide a method and a system for analyzing and learning patterns of behavior based on streams of data collected from various data sources in computer systems. In one embodiment, an information
20 security plug-in installed on computer systems in a network send data from a variety of sources to a behavior recognition system. For example, the plug-in may receive packet traffic, bit streams, and security logs.

[0016] In one embodiment, the behavior recognition system receives computer network data streams observed by the information security plug-in. The behavior
25 recognition system analyzes the data streams to identify recurring patterns.

[0017] To do so, in one embodiment, the behavior recognition system normalizes the data streams to numerical values (e.g., within a range of 0 to 1) and generates vectors of the normalized data. In one embodiment, the behavior recognition system generates a linguistic model from the normalized vectors. Specifically, a neuro-
30 linguistic module in the behavior recognition system clusters the normalized vectors and assigns a distinct symbol for each cluster. The symbols act as a building block

for the linguistic model. The neuro-linguistic module identifies combinations of symbols, i.e., words (nouns, adjective, verbs, and the like), based on a frequency of symbols recurring together in the data. That is, the neuro-linguistic module builds a dictionary that includes the words by identifying symbol combinations appearing in the
5 normalized data. Once the dictionary is mature (e.g., after identifying a specified amount of statistically relevant words), the neuro-linguistic module identifies combinations of words from the dictionary, i.e., phrases, based on a frequency of words in the dictionary which occur together in the data. In addition, the neuro-linguistic module may also evaluate unusualness scores that indicate how infrequent
10 a given symbol, word, or phrase appears in the data relative to other symbols, words, and phrases, respectively. The neuro-linguistic module may output the symbols, words, phrases, and unusualness scores to a cognitive module in the behavior recognition system. The unusualness score (for a letter, word, or phrase observed in input data) indicates how infrequently the letter, word, or phrase has occurred relative
15 to past observations. Thus, the behavior recognition system may use the unusualness scores to measure how unusual a current syntax is relative to stable models. That is, the system generates a stable model of symbols (i.e., letters), stable model of words built from the symbols (i.e., a dictionary) and a stable model of phrases built from the words (i.e., a syntax). The behavior recognition system then
20 compares current observations of letters, words, and phrases generated from the same input data stream (e.g., normalized vectors generated from network and security data) against the stable models to measure of how unusual a current observation is relative to the sable models of letters, words, and phrases. Further, the current observations update the stable models, allowing for changes in the input
25 stream, over time, to be reflected in the models generated from the input stream. The cognitive module uses the output generated by the stable models and ongoing observations to learn and recognize behaviors.

[0018] In one embodiment, the cognitive module identifies patterns based on the data streams and builds memories to learn the expected patterns, as represented by
30 different combinations of phrases, built from words in the dictionary, which themselves are built from symbols assigned to observations of the observed data. Over time, the symbols, words, and phrases generated from the observed data “mature.” Thereafter, upon observing subsequent events in the data streams, the

cognitive module is able to detect anomalous activity and generate alerts when anomalous activity is observed.

[0019] The behavior recognition system learns normal activity occurring over a computer network based on data provided by the plug-ins installed in networked systems. As a result, the behavior recognition system may detect network activity that is abnormal based on what has been observed as normal activity, without relying on training data to do so. Advantageously, this approach avoids the need for pre-defined attack signatures, or any predefined patterns, to detect suspicious behavior to systems and applications. That is, the behavior recognition system teaches itself an alphabet, a dictionary, and a grammar of phrases. Once taught, the behavior recognition system learns patterns of phrases and alerts when an observed pattern is sufficiently distinct from prior observation.

[0020] Note, the following description relies on specific examples of network and information security data provided as input to a behavior recognition system. However, one of skill in the art will recognize that embodiments are applicable in other contexts related to observing incoming network security data and analyzing the observed data to detect anomalous behavior.

[0021] Figure 1 illustrates an example computing environment 100, according to one embodiment. As shown, the computing environment 100 includes one or more computer systems 105_{1-N} connected via a network 110. Further, the computing environment includes a behavior recognition system 115. In one embodiment, the computing environment 100 is representative of an enterprise network. A computer system 105 may be a physical computing system (e.g., a desktop computer, a laptop computer, etc.) or may be a virtual computing instance executing within a computing cloud.

[0022] In one embodiment, each of the computer systems 105_{1-N} includes an information security plug-in 108_{1-N} . An information security plug-in 108 is configured to receive (or obtain) data from a variety of sources of a computer system 105. Each data source is associated with an identifier. Thus, when an information security plug-in 108 receives data from a particular source, the information security plug-in 108 may receive an identifier, timestamp information, and raw numerical values representing the actual data. Examples of the data sources include packet traffic, datagram traffic, Simple Network Management Protocol (SNMP) trap data, and the like. For instance,

the information security plug-in 108 may receive information raw numerical data associated with incoming and outgoing packet traffic, such as source addresses, destination addresses, etc. Other examples may include information related to disk mounts and physical accesses. For instance, if an individual inserts a flash drive into a USB port of a computer system 105 or mounts an external hard disk drive to the system, the information security plug-in 108 may receive a stream of data corresponding to the event (e.g., as raw numbers and identifiers associated with the flash drive, USB port, etc.).

[0023] In one embodiment, the information security plug-in 108 sends raw data to a driver in the behavior recognition system 115 that normalizes the data. The driver also evaluates statistics associated with each data value, such as a recent average and historical high and low values. Normalizing the data input by the information security plug-in 108 provides a standard format that allows the behavior recognition system 115 to analyze the data and learn patterns of behavior by observing the data streams of multiple information security plug-ins monitoring the computer network.

[0024] Illustratively, the behavior recognition system 115 includes one or more central processing units (CPUs) 120, one or more graphics processing units (GPUs) 121, network and I/O interfaces 122, a storage 124 (e.g., a disk drive, optical disk drive, and the like), and a memory 125 that includes a sensor management module 130, a sensory memory component 135, and a machine learning engine 140. In one embodiment, the behavior recognition system 115 may include one or more physical computing systems connected via a network (e.g., a local area network, the Internet, etc.). In another embodiment, the behavior recognition system 115 may be cloud-computing resources connected over the network.

[0025] The CPU 120 retrieves and executes programming instructions stored in the memory 125 as well as stores and retrieves application data residing in the storage 124. In one embodiment, the GPU 121 implements a Compute Unified Device Architecture (CUDA) developed by NVIDIA. Further, the GPU 121 provides general purpose processing using the parallel throughput architecture of the GPU 121 to more efficiently retrieve and execute programming instructions stored in the memory 125 and also to store and retrieve application data in the storage 124. Taking advantage of the parallel computing elements of the GPU 121 allows the behavior recognition system 115 to process incoming data from multiple sources

(e.g., the computer systems 105_{1-N}). As a result, the behavior recognition system 115 may scale to evaluate a large number of information streams.

5 [0026] The sensor management module 130 provides an information security driver that receives data from each information security plug-in 108. The information security driver may receive the data at specified intervals (e.g., once a minute, once every thirty minutes, once every thirty seconds, etc.). As stated, the information security driver normalizes the data sent by each of the information security plug-ins 108. In one embodiment, the sensor management module 130 may enable or disable data sources to be monitored by information security plug-ins 108 (e.g., in response to a request sent by a management console of the behavior recognition system 115).
10 For example, if the management console requests the behavior recognition system 115 to monitor outgoing traffic from a particular computer system in the network, the sensor management module 130 determines the information security plug-in 108 associated with the computer system and enables outgoing traffic monitoring for that
15 plug-in.

[0027] The sensory memory component 135 is a data store that transfers large volumes of data from the sensor management module 130 to the machine learning engine 140. The sensory memory component 135 stores the data as records. Each record may include an identifier, a timestamp, and a data payload. Further, the
20 sensory memory component 135 aggregates incoming data in a time-sorted fashion. Storing incoming data from each of the information security plug-ins 108 in a single location where the data may be aggregated allows the machine learning engine 140 to process the data efficiently. Further, the behavior recognition system 115 may reference data stored in the sensory memory component 135 in generating alerts for
25 anomalous activity. In one embodiment, the sensory memory component 135 may be implemented in via a virtual memory file system in the memory 123. In another embodiment, the sensory memory component 135 is implemented using a key-value share.

[0028] The machine learning engine 140 receives data output from the sensor management module 135. Generally, components of the machine learning engine 140 generate a linguistic representation of the normalized vectors. As described further below, to do so, the machine learning engine 140 clusters normalized values having similar features and assigns a distinct symbol to each cluster, The machine
30

learning engine 140 may then identify recurring combinations of symbols (i.e., words) in the data. The machine learning engine 140 then similarly identifies recurring combinations of words (i.e., phrases) in the data.

5 [0029] The resulting linguistic model provides grammatical statements that the machine learning engine 140 uses to describe events occurring in the data. The machine learning engine 140 analyzes the linguistic model to learn patterns of behavior. The machine learning engine 140 may store memories of learned patterns and determine statistical frequencies of patterns recurring in incoming data. Thus, over time, the machine learning engine 140 learns patterns used to determine
10 whether a current observation is a normal or abnormal event. That is, rather than rely on pre-defined attack signatures and activities defined in advance, the machine learning engine 140 maintains memories of what has been observed in the linguistic representation of the data to determine expected behavior. Thereafter, the machine learning engine 140 may decide whether the behavior of an observed event is
15 anomalous based on prior learning.

[0030] Data describing whether a normal or abnormal event has occurred or what the event is can be provided to an output device to issue alerts. Such an output device may present an alert message on a graphical user interface screen on a management console.

20 [0031] Figure 2 further illustrates the information security plug-in 108, according to one embodiment.

[0032] As shown, examples of such data 205 may include security logs, disk mount information, physical access information, disk I/O logs, IP address logs, security tool data (e.g., WireShark alerts, packet sniffer alerts, etc.), memory usage
25 information, routing information, packet traffic, datagram traffic, segment traffic, port traffic, and SNMP trap data. Of course, the data 205 are merely examples of data the information security plug-in 108 may send to the machine learning engine 140. Further, the information security plug-in 108 may include multiple sensor modules corresponding to each type of data 205 to receive. For example, one such sensor
30 module may receive octets of IP address information of incoming and outgoing network data. Another sensor module may receive data related to an external hard drive being mounted to a given computer system. The data may include port information, disk size, file system information, etc.

[0033] In one embodiment, the information security plug-in 108 sends the data 205 to an information security driver 210 within the sensor management module 130 hosted on the behavior recognition system 115. The information security driver 210 may include a data collection component 211 and a normalization component 212.

5 The data collection component 211 receives data 205 from the information security plug-in 108. The information security plug-in 108 may send the data to the data collection component 211 at an episodic rate, e.g., once every five seconds, once every ten seconds, etc. In one embodiment, the normalization component 212 converts raw data sent by an information security plug-in 108, e.g., to a numerical

10 value ranging from 0 to 1. In addition, the normalization component 212 evaluates additional statistics to associate with each value, such as a recent average and historical high and low values. The information security driver 210 may use such values to readjust normalization techniques in the event the information security driver 210 is modified. Further, normalizing the information security data provides the

15 behavior recognition system 115 with a standard format for analyzing information corresponding to data source 205. For example, assume the information security plug-in 108 sends data corresponding to an IP address to the data collection component 211. The normalization component 212 converts each octet of the IP address to a numerical value between 0 and 1. The behavior recognition system 115

20 may determine, from the normalized data, IP address locations that are regularly (or not regularly) accessed by a computer system in the network. Similarly, if the information security plug-in 108 sends data from security logs to the information security driver 210, the normalization component 212 converts the data to normalized values between 0 and 1. Thus, regardless of the information security plug-ins 108

25 used to monitor network communications or computer security data, or the form or presentation of that information, the data itself is standardized to a common format, e.g., a vector with a number of fields including a name, a timestamp, an observation ID, and a collection of floating point values, one for each normalized observation.

[0034] Figure 3 illustrates a system architecture of the behavior recognition system 115, according to one embodiment. As shown, the sensor management module 130 and the machine learning engine 140 communicate via a persistence layer 310.

[0035] The persistence layer 310 includes multiple data stores that maintain information used by components of the behavior recognition system 115. For example, the persistence layer 310 includes data stores that maintain information

describing properties of the information security plug-ins 108_{1-N} sending data to the behavior recognition system 115, system properties (e.g., serial numbers, available memory, available capacity, etc. of the behavior recognition system 115), and properties of the information security driver 210 (e.g., active plug-ins 118, active sensors associated with each plug-in 118, normalization settings, etc.). Other data stores may maintain learning model information, system events, and behavioral alerts. In addition, the sensory memory component 135 resides in the persistence layer 310.

[0036] The machine learning engine 140 itself includes a neuro-linguistic module 315 and a cognitive module 325. The neuro-linguistic module 315 performs neural network-based linguistic analysis of normalized input data to describe activity observed in the data. However, rather than describing the activity based on pre-defined objects and actions, the neuro-linguistic module 315 develops a custom language based on symbols, e.g., letters, generated from the input data. As shown, the neuro-linguistic module 315 includes a data transactional memory (DTM) component 316, a classification analyzer component 317, a mapper component 318, a lexical analyzer component 319, and a perceptual associative memory (PAM) component 320.

[0037] In one embodiment, the DTM component 316 retrieves the normalized vectors of the network security streams from the sensory memory component 135 and stages the network security data streams in the pipeline architecture provided by the GPU 121.

[0038] The classification analyzer component 317 evaluates the normalized network security streams organized by the DTM component 505 and maps the data on a neural network. In one embodiment, the neural network is a combination of a self-organizing map (SOM) and an adaptive resonance theory (ART) network. The neural network clusters the network security streams based on values occurring repeatedly in association with each other.

[0039] The mapper component 318 identifies symbols (i.e., builds an alphabet of letters) based on the clustered network security streams. Specifically, the mapper component 318 determines a statistical distribution of the clustered data and assigns a symbol to data belonging to a same cluster. In one embodiment, a symbol provides

a "fuzzy" (in the artificial intelligence sense of the word) representation of the data belonging to a given cluster. The symbol may be described as a letter of an alphabet.

[0040] Further, the mapper component 318 is adaptive. That is, the mapper component 318 may identify new symbols corresponding to new clusters generated from the normalized data, as such clusters are reinforced to be statistically relevant. The mapper component 318 learns on-line and may merge similar observations to a more generalized symbol. Specifically, when the mapper component 318 generates a cluster, the mapper component 318 assigns a symbol to that cluster. Thereafter, as additional observations map to that cluster are observed, the corresponding cluster may "mature," i.e., reach a level of statistical significance. Once a given cluster matures, the mapper component 318 begins sending that symbol to the lexical analyzer component 319 in response to normalized data that matches the cluster. In one embodiment, the mapping component 318 limits the set of symbols that can be sent to the lexical component 319 to the most statistically relevant clusters. In practice, outputting symbols (i.e., letters) assigned to the top thirty-two clusters has shown to be effective. However, other amounts may also prove effective, such as the top sixty-four or 128 most frequently recurring symbols. Note, over time, the most frequently observed symbols may change as different clusters increase in statistical significance.

[0041] In one embodiment, the mapper component 318 may decay symbols over time for particular data that the mapper component 318 no longer observes in the normalized vectors.

[0042] In one embodiment, the mapper component 318 evaluates an unusualness score for each symbol. The unusualness score is based on the frequency of a given symbol relative to other symbols observed. The unusualness score may increase or decrease over time as the neuro-linguistic module 315 receives additional data.

[0043] The mapper component 318 sends a stream of the symbols (e.g., letters), timestamp data, unusualness scores, and statistical data (e.g., a representation of the cluster associated with a given symbol) to the lexical analyzer component 319. The lexical analyzer component 319 builds a dictionary based on symbols output from the mapper component 318. In practice, the mapper component 318 may need approximately 5000 observations (i.e., normalized vectors of input data) to generate a stable alphabet of symbols.

[0044] The lexical analyzer component 319 builds a dictionary that includes combinations of symbols, e.g., words, from the symbols transmitted by the mapper component 318. The lexical analyzer component 319 identifies repeating patterns and sub-patterns of letters occurring in the letters output from the mapper component 318 and calculates frequencies of the patterns occurring throughout the symbol stream. The combinations of symbols may represent a particular activity, event, etc.

[0045] In one embodiment, the lexical analyzer component 319 limits the length of words in the dictionary to allow the lexical analyzer component 319 to identify a number of possible combinations without adversely affecting the performance of the behavior recognition system 115. Further, the lexical analyzer component 319 uses a sliding window based on the maximum length to identify words for the dictionary. For example, assume the maximum length of a word may be five symbols. In this case, the lexical analyzer component 319 may use a sliding window of ten symbols to identify words (and sub-words within the words) within the window. In practice, limiting a word to a maximum of five or six symbols has shown to be effective.

[0046] Like the mapper component 318, the lexical analyzer component 319 is adaptive. That is, the lexical analyzer component 319 may learn, decay, and generate words in the dictionary over time. Further, the lexical analyzer component 319 may determine an unusualness score for each word based on how frequently the word recurs in the data. The unusualness score may increase or decrease over time as the neuro-linguistic module 315 processes additional data.

[0047] In addition, as additional observations (i.e., symbols) are passed to the lexical analyzer component 319 and identified as a given word, the lexical analyzer component 319 may determine that the word has matured. Once a word has matured, the lexical analyzer component 319 may output observations of that word to the PAM component 320. In one embodiment, the lexical analyzer component 319 limits words sent to the PAM component 320 to the most statistically relevant words. In practice, outputting occurrences of the top 1000 most frequently occurring words has shown to be effective. Note, over time, the most frequently observed words may change as the observations of incoming letters change in frequency (or as new letters emerge by the clustering of input data by the mapper component 318).

[0048] Once the lexical analyzer component 319 has built the dictionary (i.e., identifies words that have a dynamically determined statistical relevance), the lexical

analyzer component 319 sends data that includes a stream of the symbols, words, timestamp data, unusualness scores, and statistical calculations to the PAM component 320. The PAM component 320 builds a syntax of phrases with from the words output by the lexical analyzer component 319. In practice, lexical analyzer component 319 may build a useful dictionary of words after receiving approximately 5 15,000 observations (i.e., input letters from the mapper component 318).

[0049] The PAM component 320 identifies a syntax of phrases based on the sequence of words output from the lexical analyzer component 319. Specifically, the PAM component 320 receives the words identified by the lexical analyzer component 10 319 generates a connected graph, where the nodes of the graph represent the words, and the edges represent a relationship between the words. The PAM component 320 may reinforce or decay the links based on the frequency that the words are connected with one another in a data stream.

[0050] Similar to the mapper component 318 and the lexical analyzer component 15 319, the PAM component 320 determines an unusualness score for each identified phrase based on how frequently the phrase recurs in the linguistic data. The unusualness score may increase or decrease over time as the neuro-linguistic module 315 processes additional data.

[0051] Similar to the lexical analyzer component 319, the PAM component 320 20 may limit the length of a given phrase to allow the PAM component 320 to be able to identify a number of possible combinations without adversely affecting the performance of the behavior recognition system 115. Further, PAM component 320 uses a sliding window based on the maximum length to identify phrases. For example, assume the maximum length of a word may be eight. In this case, the 25 lexical analyzer component 319 may use a sliding window of sixteen symbols to identify phrases (and sub-phrases) within the window.

[0052] The PAM component 320 identifies syntax phrases over observations of words output from the lexical analyzer component 319. As observations of words accumulate, the PAM component 320 may determine that a given phrase has 30 matured, i.e., a phrase has reached a measure of statistical relevance. The PAM component 320 then outputs observations of the that phrase to the cognitive module 325. The PAM component 320 sends data that includes a stream of the symbols, words, phrases, timestamp data, unusualness scores, and statistical calculations to

the semantic memory 330 of the cognitive module 325. In practice, the PAM component 320 may obtain a meaningful set of phrases after observing about 5000 words from the lexical analyzer component 319. The semantic memory 330 may identify complex passages of multiple phrases from the phrases output by the PAM component 320.

[0053] After maturing, the generated letters, words, and phrases form a stable neuro-linguistic model of the information security data that the behavior recognition system 115 uses to compare subsequent observations of letters, words, and phrases against the stable model. The neuro-linguistic module 315 updates the linguistic model as new data is received. Further, the neuro-linguistic module 315 may compare a currently observed syntax to the model. That is, after building a stable set of letters, the neuro-linguistic module 315 may build a stable model of words (e.g., a dictionary). In turn, the neuro-linguistic module 315 may be used to build a stable model of phrases (e.g., a syntax). Thereafter, when the neuro-linguistic module 315 receives subsequent normalized information security data, the module 315 can output an ordered stream of symbols, words, and phrases, all of which can be compared to the stable model to identify interesting patterns or detect deviations occurring in the stream of input data.

[0054] As shown, the cognitive module 325 includes a workspace 326, a semantic memory 330, codelet templates 335, episodic memory 340, long term memory 345, and an anomaly detection component 350.

[0055] In one embodiment, the workspace 326 provides a computational engine for the machine learning engine 140. For example, the workspace 326 may retrieve linguistic data having statistical relevance (e.g., phrases that occur frequently in the data) from the episodic memory 340 and the longterm memory 345, select codelet templates 335 to execute, etc. Further, the workspace 326 receives the symbols, words, and phrases (as well as unusualness scores) from the neuro-linguistic module 315 and stores these data in the semantic memory 330.

[0056] The workspace 326 retrieves the neuro-linguistic data from semantic memory 710 and disseminates this data to different portions of the cognitive module 210 as needed

[0057] The episodic memory 340 stores linguistic observations related to a particular episode in the immediate past and may encode specific details, such as the “what” and the “when” of a particular event.

5 [0058] The long-term memory 345 stores generalizations of the linguistic data with particular episodic details stripped away. In this way, when a new observation occurs, memories from the episodic memory 340 and the long-term memory 345 may be used to relate and understand a current event, i.e., the new event may be compared with past experience (as represented by previously observed linguistic data), leading to both reinforcement, decay, and adjustments to the information stored
10 in the long-term memory 345, over time. In a particular embodiment, the long-term memory 345 may be implemented as an ART network and a sparse-distributed memory data structure. Importantly, however, this approach does not require events to be defined in advance.

[0059] The codelet templates 335 provide a collection of executable codelets, or
15 small pieces of code that evaluate different sequences of events to determine how one sequence may follow (or otherwise relate to) another sequence. More generally, a codelet may detect interesting patterns from the linguistic representation of information security data. For instance, a codelet may compare a current observation (i.e., a current phrase instance with what has been observed in the past) with
20 previously observed activity stored in the semantic memory 330. By repeatedly scheduling codelets for execution, copying memories and percepts to/from the workspace 326, the cognitive module 325 performs a cognitive cycle used to observe, and learn, about patterns of behavior that occur within the linguistic data.

[0060] The anomaly detection component 350 evaluates unusualness scores sent
25 by the neuro-linguistic module 315 to determine whether to issue an alert in response to some abnormal activity indicated by the unusualness scores. Specifically, the anomaly detection component 350 provides probabilistic histogram models (e.g., an unusual lexicon model, an unusual syntax model, and an unusual model) which represent the unusualness scores. The anomaly detection component 350 may send
30 alert data to an output device, where an administrator may view the alert, e.g., via a management console.

[0061] The cognitive module 325 performs learning analysis on the linguistic content delivered to semantic memory 330 (i.e., the identified symbols, words,

phrases) by comparing new observations to the learned patterns kept in semantic memory 330 and then estimating the rareness of these new observations.

[0062] Specifically, the anomaly component 725 evaluates the unusualness scores of each of the symbols, words, and phrases to identify abnormal occurrences in the observed data. Once an anomalous observation has been identified, the anomaly component may issue an alert (e.g., notify an administrator or user of the behavior recognition system 115).

[0063] Figure 4 illustrates a method 400 for normalizing data sent by an information security plug-in to the information security driver 210 in the behavior recognition system 115, according to one embodiment.

[0064] At step 405, the information security driver 210 receives data 205 from the information security plug-in 108. The information security plug-in 108 sends data to the information security driver 210 at pre-determined intervals (e.g., once every three seconds). The data may originate from several sources, such as IP address tables, security tool alerts, access logs, etc.

[0065] At step 410, the information security driver 210 normalizes raw data to a numerical value within a range from 0 to 1 relative to the source type. For example, values associated with IP address information are normalized from 0 to 1 relative to other IP address information values.

[0066] After normalizing the values, at step 415, the information security driver 210 evaluates additional statistics to associate with each value, such as a recent average and historical high and low values for each of the values. The information security driver 210 may use such values to readjust normalization techniques in the event the information security plug-in 108 is modified.

[0067] At step 420, the information security driver 210 sends a vector containing the normalized information security data to the sensory memory component 135. As stated, the sensory memory component 135 stores the normalized values and associated data. The neuro-linguistic module 205 may then retrieve the normalized values from the sensory memory component 135 and generate a linguistic model of the normalized information security data thereafter.

[0068] Figure 5 illustrates a method 500 for generating a linguistic model of normalized sensor data, according to one embodiment. At step 505, the DTM

component 316 receives the normalized vectors of traffic of an enterprise network from sensory memory 135.

[0069] At step 510, the classification analyzer component 317 clusters the normalized values based on values recurring in association with one another. In addition, the mapper component 318 generates symbols to associate with each cluster. The mapper component 318 also evaluates the frequency that each symbol occurs in the ordered input data stream. Thereafter, the mapper component 318 determines a unusualness score for each symbol. The mapper component 318 continues to identify symbols and evaluate statistical frequencies of the identified symbols until the mapper component 318 reaches a mature state (e.g., identified a certain number of symbols), so that the lexical analyzer component 319 may reliably identify a lexicon, i.e., a dictionary, based on the symbols. In practice, a minimum threshold of approximately a hundred thousand symbols has shown to be effective.

[0070] At step 515, the lexical analyzer component 319 generates words, i.e., combinations of symbols. To do so, the lexical analyzer component 319 identifies repeating patterns and sub-patterns of symbols and calculates frequencies at which each pattern occurs throughout the ordered input symbol stream. Further, the lexical analyzer component 319 determines an unusualness score for each word based on the calculated frequencies. The lexical analyzer component 319 continues to identify words and evaluate statistical frequencies of the identified words until reaching a mature state (e.g., a minimum amount of statistically relevant words have been identified for the dictionary).

[0071] At step 520, the PAM component 320 generates a percept (syntax) of the identified words based on probabilistic relationships of the words. Further, the PAM component 320 calculates frequencies of phrases occurring in the data. Doing so allows the PAM component 320 to determine an unusualness score for each phrase based on the frequencies. The PAM component 320 continues to identify words and evaluate statistical frequencies of the identified words until reaching a mature state (e.g., a specified amount of data has been observed).

[0072] At step 525, the PAM component 320 sends the generated symbol, word, and phrase data, along with the respective unusualness scores to the cognitive module 325. The cognitive module 325 performs learning analysis on the linguistic content delivered to semantic memory 330 (i.e., the identified symbols, words,

phrases) by comparing new observations to the learned patterns kept in semantic memory 330 and then estimating the rareness of these new observations.

[0073] Further, the anomaly detection component 350 evaluates the unusualness scores of each of the symbols, words, and phrases to identify abnormal occurrences in the observed data. Once an anomalous observation has been identified, the anomaly detection component 350 may issue an alert (e.g., notify an administrator or user of the behavior recognition system 115).

[0074] One embodiment of the present disclosure is implemented as a program product for use with a computer system. The program(s) of the program product defines functions of the embodiments (including the methods described herein) and can be contained on a variety of computer-readable storage media. Examples of computer-readable storage media include (i) non-writable storage media (e.g., read-only memory devices within a computer such as CD-ROM or DVD-ROM disks readable by an optical media drive) on which information is permanently stored; (ii) writable storage media (e.g., floppy disks within a diskette drive or hard-disk drive) on which alterable information is stored. Such computer-readable storage media, when carrying computer-readable instructions that direct the functions of the present disclosure, are embodiments of the present disclosure. Other examples media include communications media through which information is conveyed to a computer, such as through a computer or telephone network, including wireless communications networks.

[0075] In general, the routines executed to implement the embodiments of the present disclosure may be part of an operating system or a specific application, component, program, module, object, or sequence of instructions. The computer program of the present disclosure is comprised typically of a multitude of instructions that will be translated by the native computer into a machine-readable format and hence executable instructions. Also, programs are comprised of variables and data structures that either reside locally to the program or are found in memory or on storage devices. In addition, various programs described herein may be identified based upon the application for which they are implemented in a specific embodiment of the disclosure. However, it should be appreciated that any particular program nomenclature that follows is used merely for convenience, and thus the present

disclosure should not be limited to use solely in any specific application identified and/or implied by such nomenclature.

[0076] While the foregoing is directed to embodiments of the present disclosure, other and further embodiments of the disclosure may be devised without departing
5 from the basic scope thereof, and the scope thereof is determined by the claims that follow.

[0077] In sum, embodiments presented herein disclose techniques for analyzing
an ordered input stream of data sent by one or more networked computer systems to
a behavior recognition system. Once deployed on a computer network, the behavior
10 recognition system identifies and learns patterns of behavior. As a result, the
behavior recognition system, after time, is able to detect anomalous behavior
occurring within the network based on previously observed activities. Therefore,
advantageously, the behavior recognition system can readily notify a system
administrator to attacks and other suspicious activity without having to configure the
15 system with pre-defined attack signatures and the like.

[0078] While the foregoing is directed to embodiments of the present disclosure, other and further embodiments of the disclosure may be devised without departing
from the basic scope thereof, and the scope thereof is determined by the claims that
follow.

What is claimed is:

1. A method for processing streams of information security data from one or more networked computer systems, the method comprising:
 - receiving an ordered stream of normalized vectors corresponding to
 - 5 information security data obtained from one or more sensors monitoring a computer network; and
 - generating a neuro-linguistic model of the information security data by:
 - clustering the ordered stream of vectors and assigning a letter to each
 - cluster,
 - 10 outputting an ordered sequence of letters based on a mapping of the ordered stream of normalized vectors to the clusters,
 - building a dictionary of words from of the ordered output of letters,
 - outputting an ordered stream of words based on the ordered output of letters, and
 - 15 generating a plurality of phrases based on the ordered output of words.
2. The method of claim 1, further comprising:
 - comparing a current observation of letters, words, or phrases generated from
 - subsequent normalized vectors in the ordered stream against the neuro-linguistic
 - 20 model; and
 - generating alerts when the current observation of letters, words, phrases, differs from the neuro-linguistic model of the information security data.
3. The method of claim 2, wherein generating alerts when the current observation
- 25 of letters, words, phrases, differs from the neuro-linguistic model of the information security data comprises determining, for at least a first, letter, word, or phrase in the current observation an unusualness score based on a frequency of letters words, and phrases, respectively, in the letters, words, and phrases of the neuro-linguistic model; and
- 30 generating the alerts when the unusualness score for at least one of the letters, words, or phrases in the current observation exceeds a specified threshold.
4. The method of claim 1, further comprising:

comparing a current observation of letters, words, or phrases generated from subsequent normalized vectors in the ordered stream against the neuro-linguistic model; and

5 generating an unusualness score for at least one of the letters, words, or phrases in the current observation.

5. The method of claim 1, further comprising,
dynamically updating the neuro-linguistic model from subsequent normalized vectors in the ordered stream of normalized vectors.

10

6. The method of claim 1, wherein the each value of the information security data associated the normalized vectors is normalized to a value within a range of 0 to 1, inclusive.

15

7. The method of claim 1, wherein the information security data comprises network packet traffic information, disk mount event information, or security log data.

8. The method of claim 1, wherein building the dictionary comprises building a dictionary of words up to a maximum length of letters.

20

9. A computer-readable storage medium storing instructions, which, when executed on a processor, performs an operation for processing streams of data of one or more networked computer systems, the operation comprising:

25 receiving an ordered stream of normalized vectors corresponding to information security data obtained from one or more sensors monitoring a computer network; and

generating a neuro-linguistic model of the information security data by:

clustering the ordered stream of vectors and assigning a letter to each cluster,

30

outputting an ordered sequence of letters based on a mapping of the ordered stream of normalized vectors to the clusters,

building a dictionary of words from of the ordered output of letters,

outputting an ordered stream of words based on the ordered output of letters, and

generating a plurality of phrases based on the ordered output of words.

10. The computer-readable storage medium of claim 9, wherein the operation further comprises:

5 comparing a current observation of letters, words, or phrases generated from subsequent normalized vectors in the ordered stream against the neuro-linguistic model; and

generating alerts when the current observation of letters, words, phrases, differs from the neuro-linguistic model of the information security data.

10

11. The computer-readable storage medium of claim 10, wherein generating alerts when the current observation of letters, words, phrases, differs from the neuro-linguistic model of the information security data comprises determining, for at least a first, letter, word, or phrase in the current observation an unusualness score based on a frequency of letters words, and phrases, respectively, in the letters, words, and phrases of the neuro-linguistic model; and

15

generating the alerts when the unusualness score for at least one of the letters, words, or phrases in the current observation exceeds a specified threshold.

20

12. The computer-readable storage medium of claim 9, wherein the operation further comprises:

comparing a current observation of letters, words, or phrases generated from subsequent normalized vectors in the ordered stream against the neuro-linguistic model; and

25

generating an unusualness score for at least one of the letters, words, or phrases in the current observation.

13. The computer-readable storage medium of claim 9, wherein the operation further comprises,

30

dynamically updating the neuro-linguistic model from subsequent normalized vectors in the ordered stream of normalized vectors.

14. The computer-readable storage medium of claim 9, wherein the each value of the information security data associated the normalized vectors is normalized to a value within a range of 0 to 1, inclusive.
- 5 15. The computer-readable storage medium of claim 9, wherein the information security data comprises network packet traffic information, disk mount event information, or security log data.
- 10 16. The computer-readable storage medium of claim 9, wherein building the dictionary comprises building a dictionary of words up to a maximum length of letters.
17. A system, comprising:
a processor; and
a memory storing one or more application programs configured to perform an
15 operation for processing streams of data of one or more networked computer systems, the operation comprising:
receiving an ordered stream of normalized vectors corresponding to information security data obtained from one or more sensors monitoring a computer network; and
20 generating a neuro-linguistic model of the information security data by:
clustering the ordered stream of vectors and assigning a letter to each cluster,
outputting an ordered sequence of letters based on a mapping of the ordered stream of normalized vectors to the clusters,
25 building a dictionary of words from of the ordered output of letters,
outputting an ordered stream of words based on the ordered output of letters, and
generating a plurality of phrases based on the ordered output of words.
- 30 18. The system of claim 17, wherein the operation further comprises:
comparing a current observation of letters, words, or phrases generated from subsequent normalized vectors in the ordered stream against the neuro-linguistic model; and

generating alerts when the current observation of letters, words, phrases, differs from the neuro-linguistic model of the information security data.

19. The system of claim 18, wherein generating alerts when the current
5 observation of letters, words, phrases, differs from the neuro-linguistic model of the information security data comprises determining, for at least a first, letter, word, or phrase in the current observation an unusualness score based on a frequency of letters words, and phrases, respectively, in the letters, words, and phrases of the neuro-linguistic model; and
10 generating the alerts when the unusualness score for at least one of the letters, words, or phrases in the current observation exceeds a specified threshold.
20. The system of claim 17, wherein the operation further comprises:
comparing a current observation of letters, words, or phrases generated from
15 subsequent normalized vectors in the ordered stream against the neuro-linguistic model; and
generating an unusualness score for at least one of the letters, words, or phrases in the current observation.
- 20 21. The system of claim 17, wherein the operation further comprises,
dynamically updating the neuro-linguistic model from subsequent normalized vectors in the ordered stream of normalized vectors.
22. The system of claim 17, wherein the each value of the information security
25 data associated the normalized vectors is normalized to a value within a range of 0 to 1, inclusive.
23. The system of claim 17, wherein the information security data comprises
network packet traffic information, disk mount event information, or security log data.
30
24. The system of claim 17, wherein building the dictionary comprises building a dictionary of words up to a maximum length of letters.

25. A method for processing streams of data of one or more networked computer systems, the method comprising:

receiving a stream of normalized vectors corresponding to information security data monitored over a computer network;

5 generating a plurality of symbols from the stream of normalized vectors;

generating a dictionary of words from combinations of the generated plurality of symbols based on an order by which symbols appear in the stream relative to one another;

10 generating a plurality of phrases based on relationships identified between words in the dictionary, wherein each relationship is based on an order by which each word appears relative to one another; and

identifying patterns of behavior within the information security data.




Figure 1

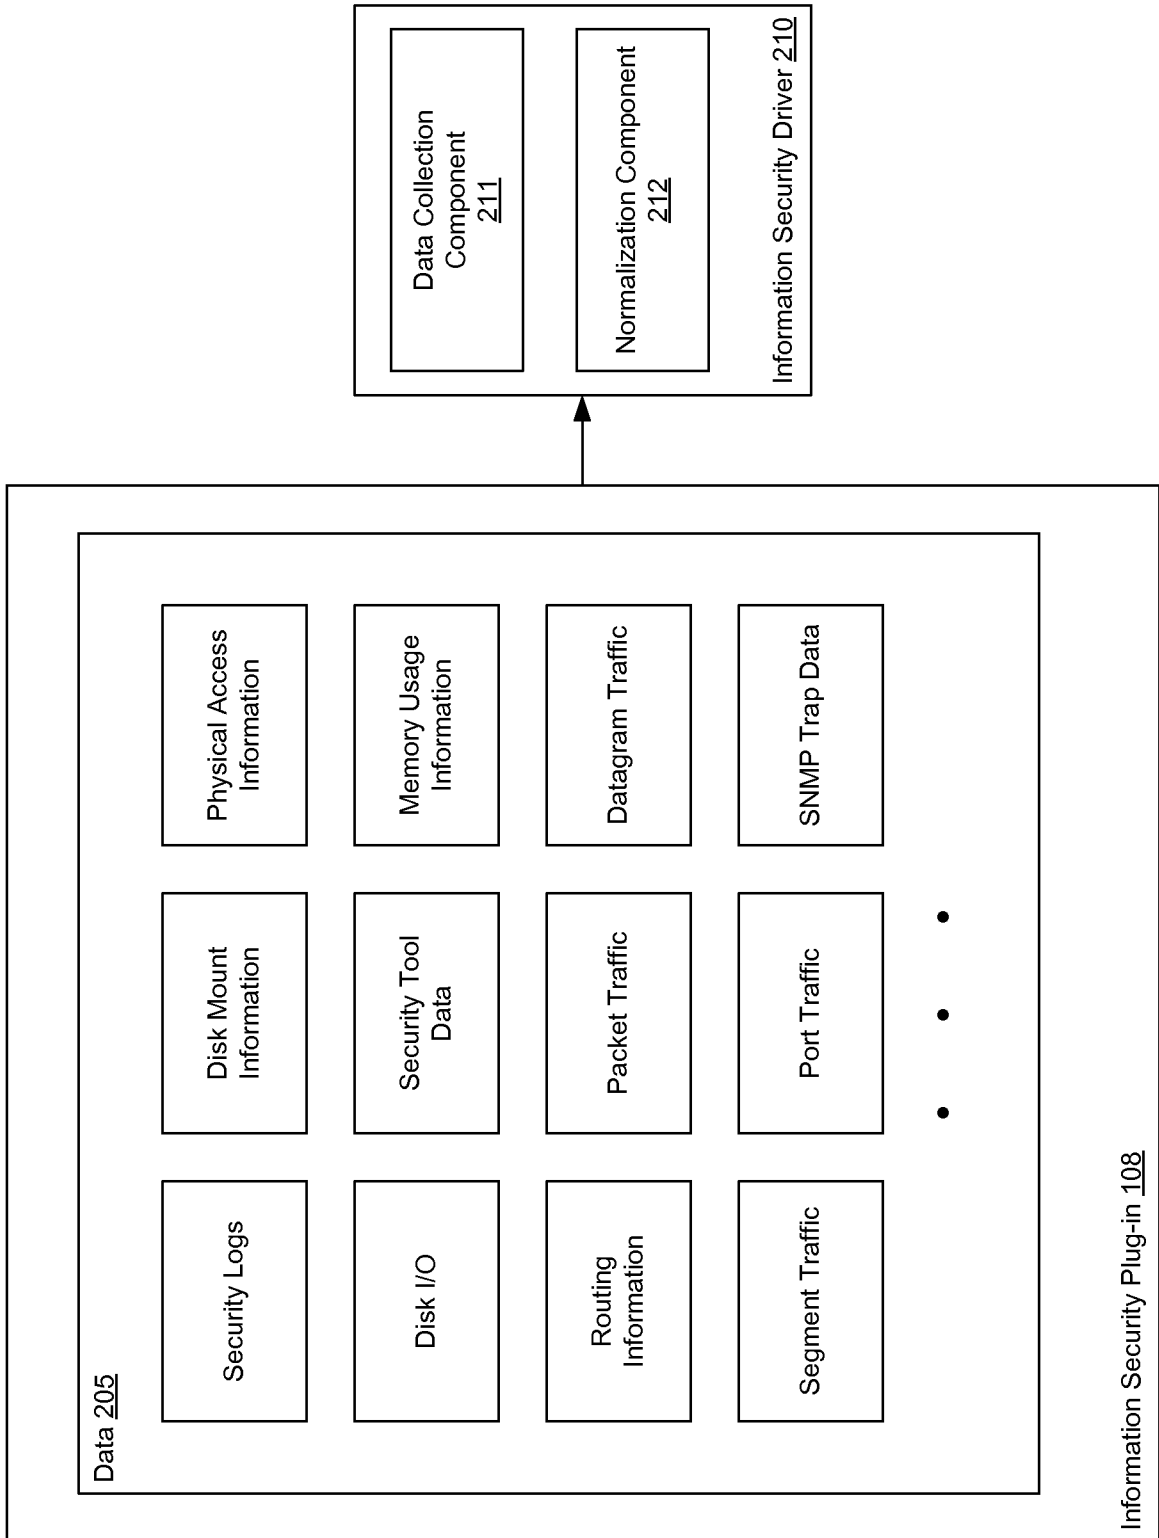


Figure 2

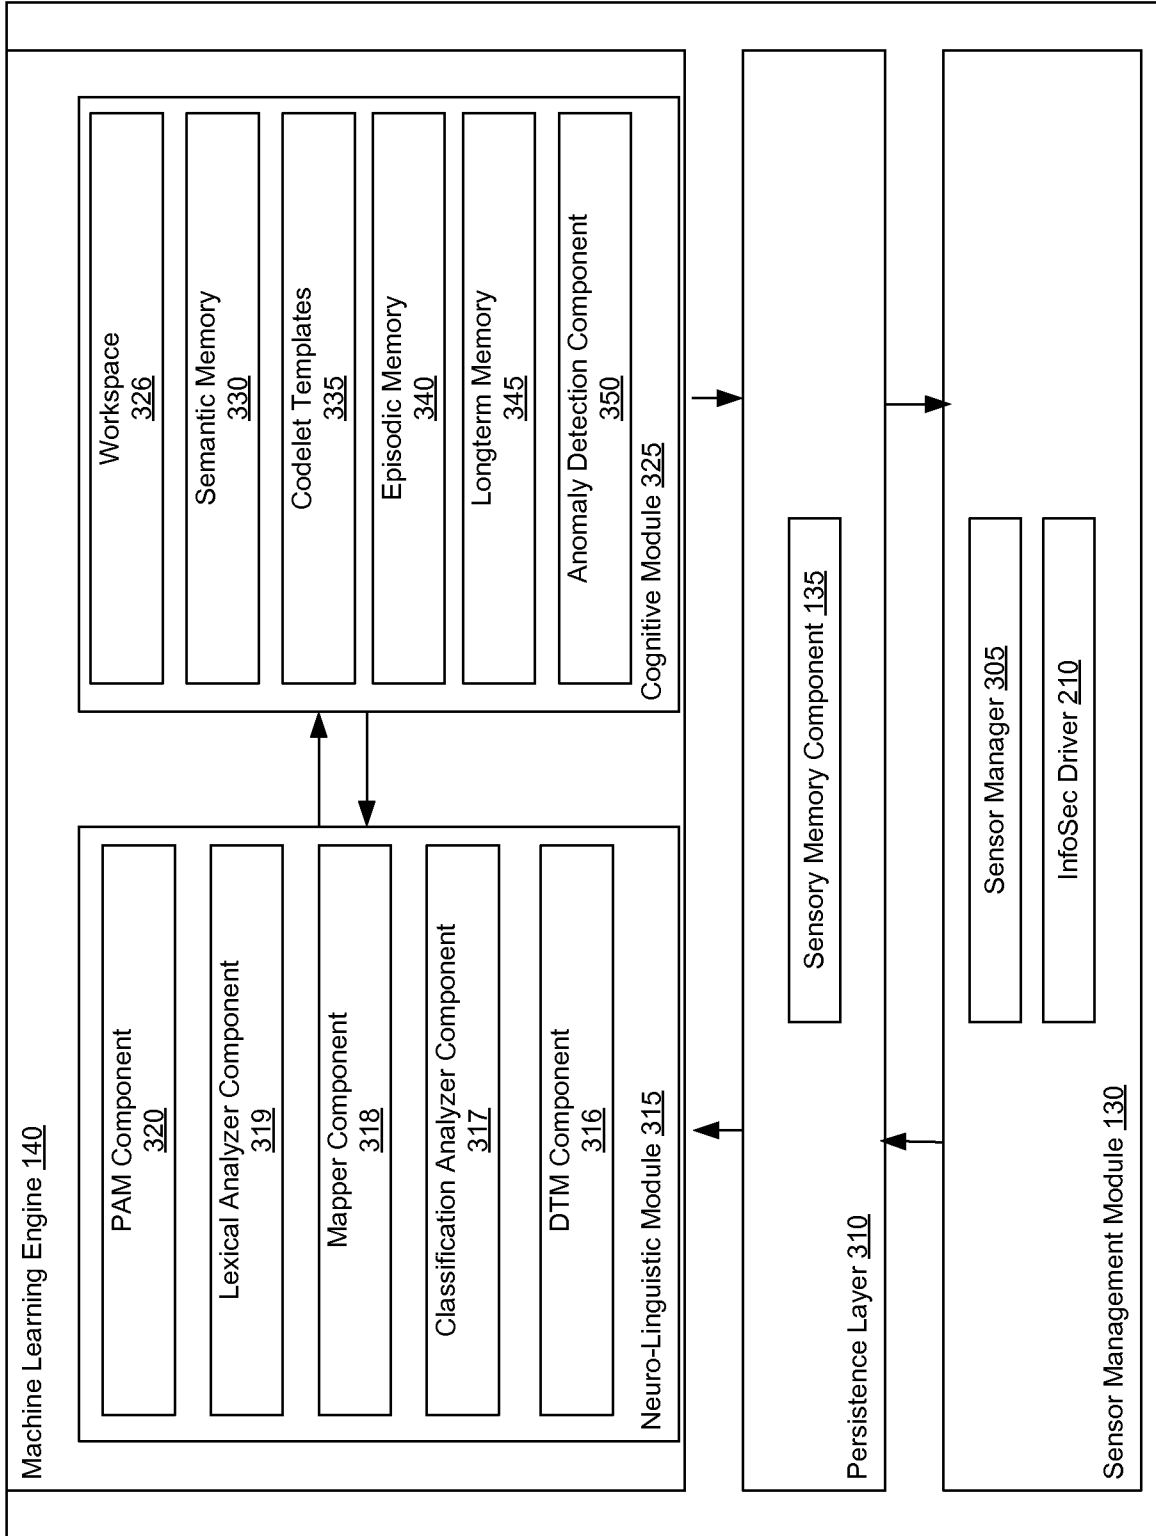


Figure 3

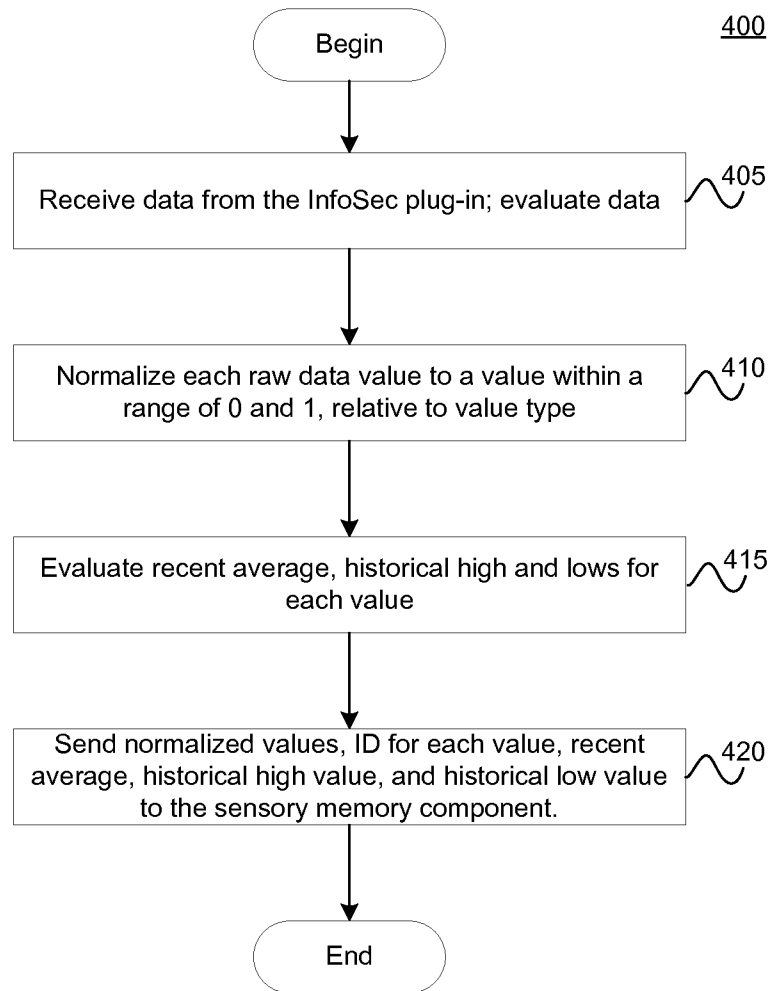


Figure 4

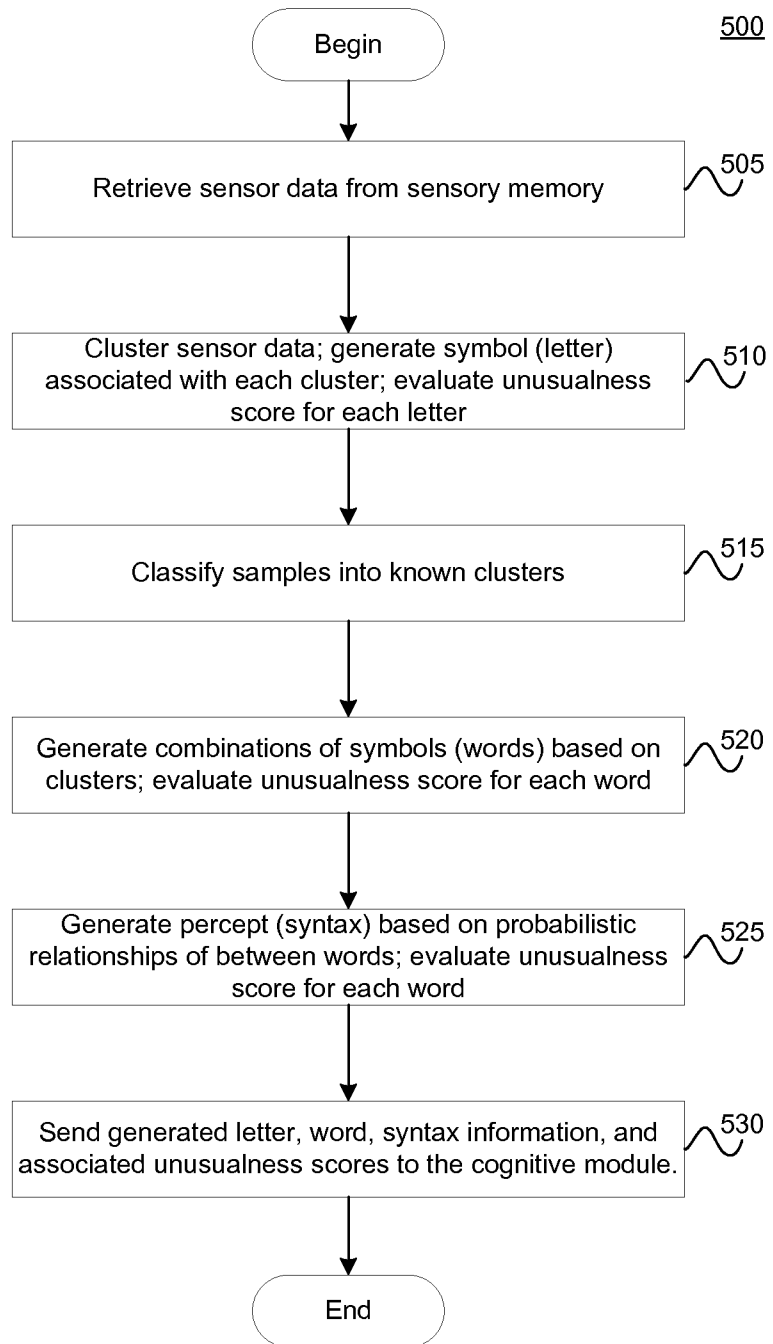


Figure 5

A. CLASSIFICATION OF SUBJECT MATTER**G06F 21/60(2013.01)i, G06N 3/02(2006.01)i**

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)

G06F 21/60; H04L 12/24; H04L 12/26; G06F 12/14; H04L 9/00; G06F 11/00; G06F 21/00; G06N 3/02

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: security, data, stream, normalized, vector, neuro-linguistic, model, letter, word, phrase

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	CN 102571486 A (SHANGHAI JIAO TONG UNIVERSITY) 11 July 2012 See paragraphs [0001], [0048]; and figure 2.	1-25
A	US 2012-0131674 A1 (THOMAS WITTENSCHLAEGER) 24 May 2012 See paragraphs [0012], [0030]-[0032]; and figures 1-2.	1-25
A	US 2006-0156404 A1 (CHRISTOPHER W. DAY) 13 July 2006 See paragraphs [0009], [0035]-[0036]; and figure 3.	1-25
A	US 2005-0262343 A1 (JIMI T. JORGENSEN et al.) 24 November 2005 See paragraphs [0022]-[0028]; and figure 1.	1-25
A	US 2008-0120720 A1 (JINHONG GUO et al.) 22 May 2008 See paragraphs [0006]-[0007], [0013]; and figures 1-2.	1-25

 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

"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

13 November 2014 (13.11.2014)

Date of mailing of the international search report

13 November 2014 (13.11.2014)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

AHN, Jeong Hwan

Telephone No. +82-42-481-8440



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2014/050604

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
CN 102571486 A	11/07/2012	CN 102571486 B	27/08/2014
US 2012-0131674 A1	24/05/2012	EP 2641100 A1 US 08683591 B2 US 2014-165201 A1 WO 2012-068443 A1 WO 2012-068443 A4	25/09/2013 25/03/2014 12/06/2014 24/05/2012 19/07/2012
US 2006-0156404 A1	13/07/2006	AU 2003-261306 A1 AU 2003-261306 A8 US 07017186 B2 US 07260846 B2 US 2004-0025044 A1 WO 2004-012063 A2 WO 2004-012063 A3	16/02/2004 16/02/2004 21/03/2006 21/08/2007 05/02/2004 05/02/2004 08/04/2004
US 2005-0262343 A1	24/11/2005	AU 2004-237046 A1 AU 2004-237046 B2 CN 1961525 A CN 1961525 B EP 1620773 A2 JP 04430666 B2 JP 2007-524892 A KR 10-0825241 B1 US 07103772 B2 US 07360087 B2 US 2004-0221163 A1 WO 2004-099940 A2 WO 2004-099940 A3	18/11/2004 28/02/2008 09/05/2007 13/04/2011 01/02/2006 10/03/2010 30/08/2007 25/04/2008 05/09/2006 15/04/2008 04/11/2004 18/11/2004 18/05/2006
US 2008-0120720 A1	22/05/2008	None	