



- (51) **International Patent Classification:**
G06F 11/34 (2006.01) *G06N 3/02* (2006.01)
G06F 11/00 (2006.01)
- (21) **International Application Number:**
PCT/US2017/026370
- (22) **International Filing Date:**
6 April 2017 (06.04.2017)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/318,867 6 April 2016 (06.04.2016) US
15/478,714 4 April 2017 (04.04.2017) US
- (71) **Applicant:** NEC LABORATORIES AMERICA, INC
[US/US]; 4 Independence Way, Suite 200, Princeton, New
Jersey 08540 (US).
- (72) **Inventors:** XU, Jianwu; 95 Oneill Court, Lawrenceville,
New Jersey 08648 (US). ZHANG, Ke; 105 Morewood
Avenue, Pittsburgh, Pennsylvania 15213 (US). ZHANG,
Hui; 69 Southfield Road, Princeton Junction, New Jersey
08550 (US). MIN, Renqiang; 7 Carlton Circle, Princeton,
New Jersey 08540 (US). JIANG, Guofei; 5 Danby Court,
Princeton, New Jersey 08540 (US).
- (74) **Agent:** KOLODKA, Joseph; NEC Laboratories America,
Inc., 4 Independence Way, Suite 200, Princeton, New Jer-
sey 08540 (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, DJ, 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, KH, KN,
KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA,
MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG,
NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS,
RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY,
TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN,
ZA, ZM, ZW.
- (84) **Designated States** (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ,
TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU,
TJ, 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:** SYSTEM FAILURE PREDICTION USING LONG SHORT-TERM MEMORY NEURAL NETWORKS




FIG. 1

(57) **Abstract:** Methods for system failure prediction include clustering log files according to structural log patterns. Feature representations of the log files are determined based on the log clusters. A likelihood of a system failure is determined based on the feature representations using a neural network. An automatic system control action is performed if the likelihood of system failure exceeds a threshold.

SYSTEM FAILURE PREDICTION USING LONG SHORT-TERM MEMORY NEURAL NETWORKS

RELATED APPLICATION INFORMATION

[0001] This application claims priority to 62/318,867 filed on April 6, 2016, incorporated herein by reference in its entirety. This application is related to an application entitled, “A MOBILE PHONE WITH SYSTEM FAILURE PREDICTION USING LONG SHORT-TERM MEMORY NEURAL NETWORKS,” attorney docket number 15113B, which is incorporated by reference herein in its entirety.

BACKGROUND

Technical Field

[0002] The present invention relates to information technology management and, more particularly, to the prediction of system failures using neural networks.

Description of the Related Art

[0003] Heterogeneous IT operational logs serve as inexpensive “sensors,” recording and indicating the health status of enterprise computer systems. Traditional system monitoring relies on manually backtracking those logs, which is both labor-intensive as well as a *post hoc* approach, i.e., identifying what went wrong after the fact. Service failure prediction becomes increasingly critical in preventing unexpected system downtime and improving the quality of service for end users.

[0004] Traditionally, the automation of this process has used text mining to extract relevant text features and patterns for failure prediction. One of the widely used features

is the term frequency – inverse document frequency (TF-IDF) feature. However, most text mining techniques based on TF-IDF lead to a high-dimensional feature space and, even more importantly, do not easily generalize to multiple log formats.

SUMMARY

[0005] A method system for failure prediction includes clustering log files according to structural log patterns. Feature representations of the log files are determined based on the log clusters. A likelihood of a system failure is determined based on the feature representations using a neural network. An automatic system control action is performed if the likelihood of system failure exceeds a threshold.

[0006] A method for system failure prediction includes clustering log files having heterogeneous log formats according to structural log patterns using unsupervised, hierarchical clustering. Feature representations of the log files are determined based on the log clusters. The feature representations consist of pattern distribution among clustered logs and term frequency – inverse document frequency. A likelihood of a system failure is determined based on the feature representations using a long short-term memory neural network a binomial distribution based on an output of the long short-term memory neural network and a binary target vector with two complementary classes. An automatic system control action is performed if the likelihood of system failure exceeds a threshold.

[0007] A system for system failure prediction includes a pattern learning module that has a processor configured to cluster a plurality of log files according to structural log patterns. A feature extraction module is configured to determine feature representations

of the log files based on the log clusters. A failure prediction module is configured to determine a likelihood of a system failure based on the feature representations using a neural network. A system control module is configured to perform an automatic system control action if the likelihood of system failure exceeds a threshold.

[0008] These and other features and advantages will become apparent from the following detailed description of illustrative embodiments thereof, which is to be read in connection with the accompanying drawings.

BRIEF DESCRIPTION OF DRAWINGS

[0009] The disclosure will provide details in the following description of preferred embodiments with reference to the following figures wherein:

[0010] FIG. 1 is a block/flow diagram of a failure prediction system in accordance with an embodiment of the present invention;

[0011] FIG. 2 is a block/flow diagram of log pattern learning in accordance with an embodiment of the present invention;

[0012] FIG. 3 is a block/flow diagram of a feature representation module in accordance with an embodiment of the present invention;

[0013] FIG. 4 is a block/flow diagram of failure prediction in accordance with an embodiment of the present invention;

[0014] FIG. 5 is a diagram of a neural network in accordance with an embodiment of the present invention;

[0015] FIG. 6 is a diagram of a neural network in accordance with an embodiment of the present invention;

[0016] FIG. 7 is a diagram of a long short-term memory cell in accordance with an embodiment of the present invention;

[0017] FIG. 8 is a block/flow diagram of a method of system failure prediction and automated system control in accordance with an embodiment of the present invention;

[0018] FIG. 9 is a block diagram of a system failure prediction and control system in accordance with an embodiment of the present invention;

[0019] FIG. 10 is a block diagram of a processing system in accordance with an embodiment of the present invention; and

[0020] FIG. 11 is a block diagram of a mobile phone device having a failure prediction and control system in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0021] In accordance with the present principles, systems and methods are provided automatically provide early warnings for potential system failure prediction by applying deep learning on heterogeneous operational logs. The present embodiments extract generic (e.g., non-textual) features from console logs that pass through a deep long short-term memory (LSTM) recurrent neural network. The LSTM recurrent neural network operates without gradient vanishing during learning.

[0022] Features are obtained based on patterns learned from unsupervised learning and are observed among similar logs, both in terms of format and content. Once clusters have been identified, the distribution of the patterns over the clusters form the features. The features are passed to the deep LSTM recurrent neural network to perform failure prediction, capturing the underlying dynamics of the system. The present embodiments

thereby provide a predictive system that achieves higher accuracy with a smaller feature set than conventional predictive systems.

[0023] Referring now to FIG. 1, a block/flow diagram of a failure prediction system 100 is shown. Given an input log file 102, log pattern learning 104 clusters the logs having similar formats together and extracts log patterns, which are stored in the global pattern set 103. It should be noted that the input log file 102 is “heterogeneous,” which is used herein to denote that the logs may be generated by systems that may have been written in different programming languages and with different logging styles, producing logs that have different structures and semantics. The log pattern learning component 104 uses unsupervised learning to automatically produce a pattern set without human input. The global pattern set 103 is periodically updated to capture the latest log patterns.

[0024] Log patterns generated by the log pattern learning component 104 are passed to a feature representation module 106. The feature representation module 106 obtains at least two types of features, including a pattern distribution among similar logs and content patterns based on term frequency – inverse document frequency (TF-IDF). These features represent the underlying dynamics and characteristics of the system being analyzed.

[0025] Failure prediction 108 takes the feature outputs from the feature representation module 106 and performs training for an LSTM neural network, upon which failure prediction 108 bases its predictions of future service failures. An LSTM recurrent neural network is able to achieve high accuracy with a small feature set as contrasted to conventional systems which may use, e.g., hundreds of features.

[0026] Referring now to FIG. 2, additional detail is provided regarding log pattern learning 104. As noted above, the present embodiments determine log format patterns based on unsupervised learning. Once the log format patterns are extracted, the input heterogeneous logs 102 are parsed using the global pattern set 103 to produce matched logs. The matched log output is provided to the feature module 106.

[0027] Application logs are generally generated from templates that are defined in the source code of the application and, thus, have a pre-defined format. Logs that come from a single source may have the same format and structure. In addition, log records can appear redundantly. This is due to the fact that some log records are used for housekeeping purposes, with the same log record being repeated many times. The existence of such regulated layouts and redundancy makes log data clustering easier than clustering in other forms of textual data. A regular layout can represent a cluster of similar logs and help reduce redundancy and summarize the meanings of log data.

[0028] Whereas conventional systems attempt to backtrack the source code to identify the layout of logged data, this is only possible when the source code is available and is therefore not generalizable to a mixture of heterogeneous logs coming from various applications, written in different programming language, with different logging styles. The present embodiments use pattern learning to capture the log structure and semantics of each field. These patterns serve as a high-level summarization of similar logs that indicate similar behaviors in the system.

[0029] Automated log clustering is performed in block 202 and uses log tokenization and unsupervised clustering. The logs are tokenized such that lower-level information from the words or phrases of each log record can be identified and retrieved. However,

heterogeneous logs from different applications and systems have different formats and, thus, different tokenizers and delimiters. Without specific domain knowledge or human inspection, there is no guarantee that particular pre-defined tokenizers will accurately tokenize the entire set of heterogeneous log data. As a result, the most general delimiter or set of delimiters should be used. The use of such a delimiter avoids interfering with potential, unknown delimiters from the log data. In one embodiment, empty space (e.g., spaces or newlines) may be used as such a general delimiter. All of the words and special symbols in a log will generally be separated by empty space.

[0030] Once the logs are tokenized into individual strings, unsupervised clustering is applied. Without domain knowledge with respect to the log formats, usage, sources, etc., a first step toward understanding and analyzing heterogeneous logs is to understand the geometric structure of the log data. Clustering serves as a way to categorize data purely based on its intrinsic properties and relations. Therefore clustering is applied on the heterogeneous logs to present the initial depiction of the data. In specific, hierarchical clustering is used to generate a hierarchical structure of the heterogeneous logs. Dense regions are searched by expanding from a certain data point toward its neighboring data points. The resulting hierarchical structure represents the data relations among its data points.

[0031] Hierarchical clustering is used not only because it can provide a coarse-to-fine view of the data, but also because the following data indexing and search is built up on a hierarchical tree structure for efficiency purposes. The results of automated log clustering in block 202 is written into log clusters 204. Logs having similar syntactic structures will

be clustered together, such that logs with similar formats will align along the same pattern.

[0032] After clustering the log data, an overall structure of all the heterogeneous logs is generated. However, patterns within each cluster are still needed to semantically understand the logs in-depth. Block 206 performs log pattern recognition to identify which field of a log record, after tokenization, is about which specific information. Since the log records have very similar formats within each cluster due, pattern recognition is performed within a cluster using sequence alignment. All the log records within a cluster are aligned together so as to identify the motifs (the most conserved and frequent common portion), and thus the common patterns, from the logs. Pairwise alignment is performed using, e.g., a Smith-Waterman process and then an Unweighted Pair Group Method with Arithmetic Mean (UPGMA) strategy is used to perform multiple alignments. The results of log pattern recognition 206 may be stored in a regular expression format and may be grouped into a log patterns file 208. The output is used to update the global pattern set 103 and serve as the final model of the input heterogeneous logs 102.

[0033] Block 210 performs log pattern matching based on the log patterns 208. The input heterogeneous logs 102 are parsed against the log patterns 208 using, e.g., a general log parsing engine. Given a regular expression format, such a tool parses the input log. Any input log will either be matched to one of the extracted log patterns 208 or will not be matched at all. If a log record is not matched, it is considered an outlier. If a log record is matched, the parsing engine automatically extracts its contents, including the contents

of individual fields such as timestamps, etc. The output of the log pattern matching 210 is provided as input to feature module 106.

[0034] Referring now to FIG. 3, additional details on the feature module 106 are shown. The feature module 106 takes the output from log pattern learning 104 and generates predictive features that can characterize the underlying dynamics of a system. Once the input logs 102 are parsed with the extracted patterns 208, each log (except for outliers) is mapped into one pattern. Thus, instead of having a collection of text logs from a time period $[t, t + \Delta t]$, a collection of patterns is extracted and the frequency p_i of each pattern i is counted. The present embodiments employ two types of features: a pattern distribution, where each element $X_{i,j}$ in the feature vector X_i is the fraction of a pattern i in the collection; and an adapted TF-IDF feature, where each pattern is an analogy to a word and the time period $[t, t + \Delta t]$ is regarded as a document.

[0035] Block 302 performs log time stamp extraction for each parsed log. Because console logs come from many different sources and may include different time stamp formats, the present embodiments standardize the time stamp formats. One exemplary format that may be used is: YYYY / MM / DD HH : MM : SS, where YYYY represents the year, MM represents the month, DD represents the day, HH represents the hour, MM represents the minute, and SS represents the second. The log time stamp extraction 302 extracts time stamps directly from the output of log pattern learning 104. If there are multiple timestamps within a given input log 102, the maximum timestamp (e.g., the latest in time) will be used to represent the log.

[0036] Block 304 generates log time series. Based on each log's represented time stamp, the present embodiments generate time series corresponding to each pattern. The

logs are discretized by a specific bin size, where the sparsity (e.g., the percentage of bins having no logs) may be kept at less than about 1%. The time granularity should be kept small to make fine-grained predictions. In one example, the time-series of pattern distributions with a ten minute resolution may have a sparsity of about 0.82%. Each time series includes histograms of each pattern over time. The higher the time series value is at a particular time stamp, the more logs correspond to the pattern at that time. Therefore the log time series captures all the dynamics of the system.

[0037] Block 306 extracts pattern distribution features as the relative distribution of each pattern in the overall logs for each time period $[t, t + \Delta t]$. Because each log pattern originated from different places in applications' source codes, the patterns provide useful information regarding the dynamics of the system for failure prediction. The pattern distribution feature may be generated in a multi-dimensional vector format. The number of dimensions equals the number of distinct patterns. Each element of the vector may be normalized so that the sum is equal to 1.

[0038] Block 308 performs a second type of feature extraction as TF-IDF. Each pattern is treated as a word and all logs occurring in each time period are treated as a single document. TF-IDF feature extraction 308 computes, for each pattern i , $\text{TF-IDF}_i = P_i \text{IDF}_i$, where P_i is the frequency of pattern i in the overall logs and IDF_i is the inverse document frequency for the pattern i , defined as the log of the total number of time periods divided by the number of time periods with the pattern i in it. The TF-IDF feature vector will have the same dimensionality as the pattern distribution feature vector produced by block 306, such that the two vectors can be combined into one and provided as an input to failure prediction 108.

[0039] Referring now to FIG. 4, additional details on failure prediction 108 are shown. Failure prediction 108 takes the output from the feature module 106 and generates results that predict the failure status of a system. An LSTM recurrent neural network 402 is a neural network designed to be better at storing and accessing historical information than classical recurrent neural networks.

[0040] The architecture of the LSTM recurrent neural network 402 will be described in greater detail below. The same architecture is used in the neural network 402 for dynamic system learning and prediction. In the learning stage, internal parameters of the neural network 402 are learned based on, e.g., back propagation through time. The input includes log features from training data and user labels on failure time periods. In the prediction stage, the neural network 402 takes log features from operational data as input and uses the output of the binary classifier 404 for failure prediction.

[0041] The present embodiments provide the probability of a system failure based on the input sequence of historical feature vectors extracted from the logs 102. If such a probability exceeds a pre-defined threshold, an early warning is reported to signal the potential system failure. An early warning signal may be very weak and thus difficult to capture by naïve models. The binary classification 404 makes use of a binary target vector d with two complementary classes. The output of the neural network 402 (expressed herein as y_t) is essentially a binary vector serving as a representation of the system status, which can be used to estimate the binomial distribution $\Pr(d_t|y_t)$. This distribution can be characterized by a softmax function at the outer layer:

$$\Pr(d_t = k|y_t) = \hat{y}_t^k = \frac{\exp(y_t^k)}{\sum_{k'=1}^K \exp(y_t^{k'})}$$

where $k = 2$ is the number of classes. For the objective function, the binary cross-entropy cost function is used based on the output at a time step t for one training example:

$$\mathcal{C} = - \sum_{k=1}^K [d_t^k \log(\hat{y}_t^k) + (1 - d_t^k) \log(\hat{y}_t^k)]$$

where the target d_t^k is either one or zero.

[0042] For the size of the neural network 402, the input size depends on the input feature dimension and output size is just 2 as a result of being a binary classification task. The layer and the number of hidden units determine the width and depth of the network. The more layers and the larger the number of hidden units, the more complex the network becomes to train.

[0043] Referring now to FIG. 5, a generalized diagram of a neural network is shown. ANNs demonstrate an ability to derive meaning from complicated or imprecise data and can be used to extract patterns and detect trends that are too complex to be detected by humans or other computer-based systems. The structure of a neural network is known generally to have input neurons 502 that provide information to one or more “hidden” neurons 504. Connections 508 between the input neurons 502 and hidden neurons 504 are weighted and these weighted inputs are then processed by the hidden neurons 504 according to some function in the hidden neurons 504, with weighted connections 508 between the layers. There may be any number of layers of hidden neurons 504, and as well as neurons that perform different functions. There exist different neural network structures as well, such as convolutional neural network, maxout network, etc. Finally, a set of output neurons 506 accepts and processes weighted input from the last set of hidden neurons 504.

[0044] This represents a “feed-forward” computation, where information propagates from input neurons 502 to the output neurons 506. Upon completion of a feed-forward computation, the output is compared to a desired output available from training data. The error relative to the training data is then processed in “feed-back” computation, where the hidden neurons 504 and input neurons 502 receive information regarding the error propagating backward from the output neurons 506. Once the backward error propagation has been completed, weight updates are performed, with the weighted connections 508 being updated to account for the received error. This represents just one variety of ANN.

[0045] Referring now to FIG. 6, an artificial neural network (ANN) architecture 600 is shown. It should be understood that the present architecture is purely exemplary and that other architectures or types of neural network may be used instead. During feed-forward operation, a set of input neurons 602 each provide an input voltage in parallel to a respective row of weights 604. The weights 604 each have a settable resistance value, such that a current output flows from the weight 604 to a respective hidden neuron 606 to represent the weighted input. The current output by a given weight is determined as $I = \frac{V}{r}$, where V is the input voltage from the input neuron 602 and r is the set resistance of the weight 604. The current from each weight adds column-wise and flows to a hidden neuron 606. A set of reference weights 607 have a fixed resistance and combine their outputs into a reference current that is provided to each of the hidden neurons 606. Because conductance values can only be positive numbers, some reference conductance is needed to encode both positive and negative values in the matrix. The currents produced by the weights 604 are continuously valued and positive, and therefore the reference weights 607 are used to provide a reference current, above which currents are

considered to have positive values and below which currents are considered to have negative values.

[0046] As an alternative to using the reference weights 607, another embodiment may use separate arrays of weights 604 to capture negative values. Each approach has advantages and disadvantages. Using the reference weights 607 is more efficient in chip area, but reference values need to be matched closely to one another. In contrast, the use of a separate array for negative values does not involve close matching as each value has a pair of weights to compare against. However, the negative weight matrix approach uses roughly twice the chip area as compared to the single reference weight column. In addition, the reference weight column generates a current that needs to be copied to each neuron for comparison, whereas a negative matrix array provides a reference value directly for each neuron. In the negative array embodiment, the weights 604 of both positive and negative arrays are updated, but this also increases signal-to-noise ratio as each weight value is a difference of two conductance values. The two embodiments provide identical functionality in encoding a negative value and those having ordinary skill in the art will be able to choose a suitable embodiment for the application at hand.

[0047] The hidden neurons 606 use the currents from the array of weights 604 and the reference weights 607 to perform some calculation. The hidden neurons 606 then output a voltage of their own to another array of weights 607. This array performs in the same way, with a column of weights 604 receiving a voltage from their respective hidden neuron 606 to produce a weighted current output that adds row-wise and is provided to the output neuron 608.

[0048] It should be understood that any number of these stages may be implemented, by interposing additional layers of arrays and hidden neurons 606. It should also be noted that some neurons may be constant neurons 609, which provide a constant voltage to the array. The constant neurons 609 can be present among the input neurons 602 and/or hidden neurons 606 and are only used during feed-forward operation.

During back propagation, the output neurons 608 provide a voltage back across the array of weights 604. The output layer compares the generated network response to training data and computes an error. The error is applied to the array as a voltage pulse, where the height and/or duration of the pulse is modulated proportional to the error value. In this example, a row of weights 604 receives a voltage from a respective output neuron 608 in parallel and converts that voltage into a current which adds column-wise to provide an input to hidden neurons 606. The hidden neurons 606 provide combine the weighted feedback signal with a derivative of its feed-forward calculation and stores an error value before outputting a feedback signal voltage to its respective column of weights 604. This back propagation travels through the entire network 600 until all hidden neurons 606 and the input neurons 602 have stored an error value.

[0049] During weight updates, the input neurons 602 and hidden neurons 606 apply a first weight update voltage forward and the output neurons 608 and hidden neurons 606 apply a second weight update voltage backward through the network 600. The combinations of these voltages create a state change within each weight 604, causing the weight 604 to take on a new resistance value. In this manner the weights 604 can be trained to adapt the neural network 600 to errors in its processing. It should be noted that

the three modes of operation, feed forward, back propagation, and weight update, do not overlap with one another.

[0050] Referring now to FIG. 7, an alternative view of an LSTM memory cell 700 is shown. The memory cell 700 takes as input a vector $\mathbf{x}$ which represents the input feature vector sequence $\mathbf{x} = (x_{t-L+1}, \dots, x_t)$ with a sequence length of L^2 . The input vector $\mathbf{x}$ is passed to a stack of multiple recurrently connected hidden layers through weighted connections to compute hidden vector sequences $\mathbf{h} = (h_{t-L+1}, \dots, h_t)$ and, finally, an output vector sequence $\mathbf{y} = (y_{t-L+1}, \dots, y_t)$. The output vector $\mathbf{y}_t$ can be used to parameterize the probability distribution $Pr(d_t|y_t)$ of the target d_t .

[0051] In particular, the memory cell includes an input gate 702, a hidden gate 704, an output gate 708, and a forget gate 710, each of which is associated with a different function over the weights 604 of the neural network 600. Each gate accepts the input vector $\mathbf{x}$ and performs a respective calculation on said input as defined below, with the input gate 702 calculating i_t , the hidden gate 704 calculating z_t , the output gate 708 calculating o_t , and the forget gate calculating f_t . A cell 706 performs a further calculation to generate c_t . An arctan function 712 operates on the output of the cell 706. These calculations are set forth below as:

$$\left\{ \begin{array}{l} i_t = \sigma(W_{xi}x_t + W_{hi}h_{t-1} + W_{ci}c_{t-1} + b_i) \\ o_t = \sigma(W_{xo}x_t + W_{ho}h_{t-1} + W_{co}c_{t-1} + b_o) \\ f_t = \sigma(W_{xf}x_t + W_{hf}h_{t-1} + W_{cf}c_{t-1} + b_f) \\ \quad z_t = W_{xc}x_t + W_{cf}h_{t-1} + b_c \\ \quad c_t = f_t c_{t-1} + i_t \tanh(z_t) \\ \quad h_t = o_t \tanh(c_t) \end{array} \right.$$

where σ is the logistic sigmoid function, b_* are the bias terms, and i, f, o , and c are the input gate, forget gate, output gate, and cell vectors respectively, all of which have the

same size as the hidden state vector h . The weight matrix W indicates the connections between gates, the cell, input and hidden states. For instance, W_{hi} is the hidden-input gate matrix, while W_{xf} is the input-forget gate matrix. The weight matrices from the cell to gate vectors (e.g., W_{ci}) are diagonal, so the m^{th} element in each gate vector only receives input from the m^{th} of the cell vector.

[0052] Referring now to FIG. 8, a method of system failure prediction is shown. Block 802 gathers the heterogeneous system logs 102 from any number of different system applications and services. As noted above, the system logs 102 may be in any format and may, therefore, be very different from one another in how they store relevant system information. To accommodate these different log formats, block 804 learns log patterns. As described above, it is specifically contemplated that block 804 may employ hierarchical clustering to group the log records into clusters of similar structure.

[0053] Once the logs are clustered, block 806 determines feature representations of the logs. Block 806 recognizes patterns within the logs and generates, e.g., two distinct features to represent the log information: TF-IDF features and a pattern distribution. Based on these features, the LSTM recurrent neural network 402 and binary classifier 404 are used to predict system failures in block 808.

[0054] Block 810 performs automatic system control actions based on the predicted system failures. These automatic controls can include any of a variety of actions, ranging from issuing an alert to a human system operator to adjusting the operation of hardware and software systems to prevent or mitigate the system failure.

[0055] Embodiments described herein may be entirely hardware, entirely software or including both hardware and software elements. In a preferred embodiment, the present

invention is implemented in software, which includes but is not limited to firmware, resident software, microcode, etc.

[0056] Embodiments may include a computer program product accessible from a computer-usable or computer-readable medium providing program code for use by or in connection with a computer or any instruction execution system. A computer-usable or computer readable medium may include any apparatus that stores, communicates, propagates, or transports the program for use by or in connection with the instruction execution system, apparatus, or device. The medium can be magnetic, optical, electronic, electromagnetic, infrared, or semiconductor system (or apparatus or device) or a propagation medium. The medium may include a computer-readable storage medium such as a semiconductor or solid state memory, magnetic tape, a removable computer diskette, a random access memory (RAM), a read-only memory (ROM), a rigid magnetic disk and an optical disk, etc.

[0057] Each computer program may be tangibly stored in a machine-readable storage media or device (e.g., program memory or magnetic disk) readable by a general or special purpose programmable computer, for configuring and controlling operation of a computer when the storage media or device is read by the computer to perform the procedures described herein. The inventive system may also be considered to be embodied in a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer to operate in a specific and predefined manner to perform the functions described herein.

[0058] A data processing system suitable for storing and/or executing program code may include at least one processor coupled directly or indirectly to memory elements

through a system bus. The memory elements can include local memory employed during actual execution of the program code, bulk storage, and cache memories which provide temporary storage of at least some program code to reduce the number of times code is retrieved from bulk storage during execution. Input/output or I/O devices (including but not limited to keyboards, displays, pointing devices, etc.) may be coupled to the system either directly or through intervening I/O controllers.

[0059] Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modem and Ethernet cards are just a few of the currently available types of network adapters.

[0060] Referring now to FIG. 9, a system failure prediction and control system 900 is shown. The system 900 includes a hardware processor and memory 904. The system 900 further includes an LSTM neural network 906 that may be implemented in software or that may, alternatively, be implemented using hardware devices that implement the neurons and weights. System logs 908 are stored in the memory 904 and may be collected from any appropriate source application or system. In addition, the system 900 includes one or more functional modules. In some embodiments, the functional modules may be implemented as software that is stored in memory 904 and that is executed by processor 902. In other embodiments, the functional modules may be implemented as one or more discrete hardware components using, e.g., application-specific integrated chips or field programmable gate arrays.

[0061] For example, a pattern learning module uses clustering to group the heterogeneous system logs 908 into clusters having similar structures. A feature extraction module 912 uses the clustered log information to extract a limited number of features that failure prediction module 914 uses to determine whether a system failure may occur in one or more applications or systems. The failure prediction module 914 employs the LSTM neural network to determine whether the information provided by the system logs, characterized by the extracted features, provides an indication of a system failure.

[0062] A system control module 916 performs actions based on the prediction generated by failure prediction module. If a failure is predicted, the system control module 916 may, for example, issue an alert to a system operator or may directly perform an action to mitigate or prevent the system failure. Such automatic actions may include, for example, system load balancing, cooling/environmental control, network control actions (e.g., rerouting traffic or blocking a particular connection or data stream), and system shut down.

[0063] Referring now to FIG. 10, an exemplary processing system 1000 is shown which may represent the system failure prediction and control system 900. The processing system 1000 includes at least one processor (CPU) 1004 operatively coupled to other components via a system bus 1002. A cache 1006, a Read Only Memory (ROM) 1008, a Random Access Memory (RAM) 1010, an input/output (I/O) adapter 1020, a sound adapter 1030, a network adapter 1040, a user interface adapter 1050, and a display adapter 1060, are operatively coupled to the system bus 1002.

[0064] A first storage device 1022 and a second storage device 1024 are operatively coupled to system bus 1002 by the I/O adapter 1020. The storage devices 1022 and 1024 can be any of a disk storage device (e.g., a magnetic or optical disk storage device), a solid state magnetic device, and so forth. The storage devices 1022 and 1024 can be the same type of storage device or different types of storage devices.

[0065] A speaker 1032 is operatively coupled to system bus 1002 by the sound adapter 1030. A transceiver 1042 is operatively coupled to system bus 1002 by network adapter 1040. A display device 1062 is operatively coupled to system bus 1002 by display adapter 1060.

[0066] A first user input device 1052, a second user input device 1054, and a third user input device 1056 are operatively coupled to system bus 1002 by user interface adapter 1050. The user input devices 1052, 1054, and 1056 can be any of a keyboard, a mouse, a keypad, an image capture device, a motion sensing device, a microphone, a device incorporating the functionality of at least two of the preceding devices, and so forth. Of course, other types of input devices can also be used, while maintaining the spirit of the present principles. The user input devices 1052, 1054, and 1056 can be the same type of user input device or different types of user input devices. The user input devices 1052, 1054, and 1056 are used to input and output information to and from system 1000.

[0067] Of course, the processing system 1000 may also include other elements (not shown), as readily contemplated by one of skill in the art, as well as omit certain elements. For example, various other input devices and/or output devices can be included in processing system 1000, depending upon the particular implementation of the same, as readily understood by one of ordinary skill in the art. For example, various

types of wireless and/or wired input and/or output devices can be used. Moreover, additional processors, controllers, memories, and so forth, in various configurations can also be utilized as readily appreciated by one of ordinary skill in the art. These and other variations of the processing system 1000 are readily contemplated by one of ordinary skill in the art given the teachings of the present principles provided herein.

[0068] Referring now to FIG. 11, a mobile phone device 1100 is shown. It should be understood that the present embodiments may be implemented in any appropriate end-user product—the mobile phone device 1100 is used solely to illustrate one application and should not be construed as limiting.

[0069] The mobile phone device 1100 includes a processor 1102 and a memory 1104. A radio interface 1106 provides communications with a mobile telephony network and may operate on any appropriate wireless communications standard. A distinct network interface adapter 1108 provides communications with a user's home network by any appropriate wireless communications standard such as, e.g., WIFI.

[0070] An operating system 1110 runs on the mobile phone device 1100 and provides a functional interface between the processor 1102, memory 1104, and one or more user applications 1112. A user interface 1114 provides a graphical and/or textual display that provides a user with information and that accepts inputs from the user. In particular, an alert module 1116 interfaces with the user interface 1114 to provide information to a user or service provider regarding the status of the components. The alert module 1116 may provide a visual indicator, for example by an icon or light source, a textual indicator, for example a notification or email, or an auditory indicator, for example a chime or other sound.

[0071] All of these different components may have respective detectors that gather data regarding their respective components and store the data as logs. As noted in detail above, the logs may have different formats, but the failure prediction and control system finds patterns within the logs, derives a set of representative features, and predicts failures for components within the mobile device 1100. Any user-installed application 1112 may contribute such logs, and the failure prediction and control system 900 can thereby provide adaptive system monitoring for comprehensive failure prediction by determining a likelihood of failure for each component.

[0072] The failure prediction and control system 900 can furthermore take actions to ameliorate or prevent such predicted failures. These actions may include, for example, turning off one or more components, decreasing processor performance to lower temperature or extend battery life, and managing app permissions if a user application 1112 is consuming too many system resources. The failure prediction and control system 900 may take action if the likelihood of failure exceeds an automatic control threshold.

[0073] The alert module 1116 can provide information to a user or service provider regarding the prediction. In one example, if the failure prediction and control system 900 predicts the failure of one or more components, the alert module 1116 can provide information to the user about the potential failure. In another embodiment, the alert module 1116 can provide information to a service provider or manager, for example in the case where the mobile phone device 1100 is deployed to employees. In this case, the manager can then proactively repair or replace the mobile phone device 1100. The alert module 1116 may generate an alert if the likelihood of a failure exceeds an alert

threshold. This threshold may be set at a lower likelihood than a threshold used to automatically control the system.

[0074] The foregoing is to be understood as being in every respect illustrative and exemplary, but not restrictive, and the scope of the invention disclosed herein is not to be determined from the Detailed Description, but rather from the claims as interpreted according to the full breadth permitted by the patent laws. It is to be understood that the embodiments shown and described herein are only illustrative of the principles of the present invention and that those skilled in the art may implement various modifications without departing from the scope and spirit of the invention. Those skilled in the art could implement various other feature combinations without departing from the scope and spirit of the invention. Having thus described aspects of the invention, with the details and particularity required by the patent laws, what is claimed and desired protected by Letters Patent is set forth in the appended claims.

WHAT IS CLAIMED IS:

1. A method for system failure prediction, comprising:
clustering a plurality of log files according to structural log patterns;
determining feature representations of the log files based on the log clusters;
determining a likelihood of a system failure based on the feature representations
using a neural network; and
performing an automatic system control action if the likelihood of system failure exceeds a threshold.
2. The method of claim 1, wherein determining feature representations comprises determining two, and only two, feature representations.
3. The method of claim 2, wherein the two feature representations consist of pattern distribution among clustered logs and term frequency – inverse document frequency.
4. The method of claim 1, wherein clustering the plurality of log files comprises performing unsupervised, hierarchical clustering.
5. The method of claim 1, wherein determining the likelihood of system failure comprises processing the feature representations using a long short-term memory recurrent neural network.

6. The method of claim 5, further comprising training the long short-term memory recurrent neural network from training data and user labels on failure time periods.

7. The method of claim 1, wherein determining the likelihood of a system failure comprises determining a binomial distribution based on an output of the neural network and a binary target vector with two complementary classes.

8. The method of claim 1, wherein the plurality of log files comprise a plurality of different log formats.

9. The method of claim 8, wherein the plurality of different log formats comprise different delimiter formats.

10. A method for system failure prediction, comprising:
clustering a plurality of log files having heterogeneous log formats according to structural log patterns using unsupervised, hierarchical clustering;
determining feature representations of the log files based on the log clusters, wherein the feature representations consist of pattern distribution among clustered logs and term frequency – inverse document frequency;
determining a likelihood of a system failure based on the feature representations using a long short-term memory neural network a binomial distribution based on an

output of the long short-term memory neural network and a binary target vector with two complementary classes; and

performing an automatic system control action if the likelihood of system failure exceeds a threshold.

11. A system for system failure prediction, comprising:

a pattern learning module comprising a processor configured to cluster a plurality of log files according to structural log patterns;

a feature extraction module configured to determine feature representations of the log files based on the log clusters;

a failure prediction module configured to determine a likelihood of a system failure based on the feature representations using a neural network; and

a system control module configured to perform an automatic system control action if the likelihood of system failure exceeds a threshold.

12. The system of claim 11, wherein the feature extraction module is further configured to determine two, and only two, feature representations.

13. The system of claim 12, wherein the two feature representations consist of pattern distribution among clustered logs and term frequency – inverse document frequency.

14. The system of claim 11, wherein the pattern learning module is further configured to perform unsupervised, hierarchical clustering.

15. The system of claim 11, wherein the failure prediction module is further configured to process the feature representations using a long short-term memory recurrent neural network.

16. The system of claim 11, wherein the failure prediction module is further configured to determine a binomial distribution based on an output of the neural network and a binary target vector with two complementary classes.

17. The system of claim 11, wherein the plurality of log files comprise a plurality of different log formats.

18. The system of claim 17, wherein the plurality of different log formats comprise different delimiter formats.

1/10




FIG. 1

2/10




FIG. 2

3/10




FIG. 3




FIG. 4




FIG. 5

5/10




FIG. 6




FIG. 7

7/10




FIG. 8

8/10




FIG. 9

9/10




FIG. 10

10/10




FIG. 11

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/026370**A. CLASSIFICATION OF SUBJECT MATTER****G06F 11/34(2006.01)i, G06F 11/00(2006.01)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 11/34; G06F 11/07; H04L 29/06; G06F 11/00; 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: log files, failure prediction, clustering, distribution, Term Frequency - Inverse Document Frequency (TF-IDF), threshold, and similar terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	TIM ZWIETASCH, 'Detecting anomalies in system log files using machine learning techniques,' Bachelor's Thesis of University of Stuttgart, 2014 Retrieved from < https://elib.uni-stuttgart.de/handle/11682/3471?locale=en > See sections 1.1, 3.2-3.5, and 4.1-4.2.	1-18
Y	US 8,887,286 B2 (CATAPHORA, INC.) 11 November 2014 See column 96, lines 38-44.	1-18
Y	PANKAJ MALHOTRA et al., 'Long Short Term Memory Networks for Anomaly Detection in Time Series,' In: proceedings on European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN 2015), pp. 89-94, 22-24 April 2015 Retrieved from < https://www.researchgate.net/publication/304782562_Long_Short_Term_Memory_Networks_for_Anomaly_Detection_in_Time_Series > See abstract and section 2.	5-6, 10, 15
A	US 2011-0185234 A1 (IRA COHEN et al.) 28 July 2011 See paragraphs [0032]-[0048] and figure 4.	1-18
A	US 2015-0089309 A1 (EMC CORPORATION) 26 March 2015 See paragraphs [0035]-[0037] and [0041]-[0060]; and figures 3 and 5.	1-18



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

30 June 2017 (30.06.2017)

Date of mailing of the international search report

30 June 2017 (30.06.2017)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsu-ro, Seo-gu, Daejeon, 35208, Republic of Korea



Facsimile No. +82-42-481-8578

Authorized officer

NHO, Ji Myong

Telephone No. +82-42-481-8528



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/026370

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 8887286 B2	11/11/2014	US 2012-137367 A1 US 2014-096249 A1	31/05/2012 03/04/2014
US 2011-0185234 A1	28/07/2011	US 8209567 B2	26/06/2012
US 2015-0089309 A1	26/03/2015	CN 104461842 A US 9542255 B2	25/03/2015 10/01/2017