Predicting Breakdowns in Cloud Services (with SPIKE)

Jianfeng Chen
jchen37@ncsu.edu
NC State University
Raleigh, NC, USA

Joymallya Chakraborty
jchakra@ncsu.edu
NC State University
Raleigh, NC, USA

Philip Clark
philip.clark@lexisnexis.com
LexisNexis Legal & Professional
Raleigh, NC, USA

Kevin Haverlock
kevin.haverlock@lexisnexis.com
LexisNexis Legal & Professional
Raleigh, NC, USA

Snehit Cherian
snehit.cherian@lexisnexis.com
LexisNexis Legal & Professional
Raleigh, NC, USA

Tim Menzies
timm@ieee.org
NC State University
Raleigh, NC, USA

ABSTRACT

Maintaining web-services is a mission-critical task where any downtime means loss of revenue and reputation (of being a reliable service provider). In the current competitive web services market, such a loss of reputation causes extensive loss of future revenue.

To address this issue, we developed *SPIKE*, a data mining tool which can predict upcoming service breakdowns, half an hour into the future. Such predictions let an organization alert and assemble the tiger team to address the problem (e.g. by reconfiguring cloud hardware in order to reduce the likelihood of that breakdown).

SPIKE utilizes (a) regression tree learning (with CART); (b) synthetic minority over-sampling (to handle how rare spikes are in our data); (c) hyperparameter optimization (to learn best settings for our local data) and (d) a technique we called “topology sampling” where training vectors are built from extensive details of an individual node plus summary details on all their neighbors.

In the experiments reported here, *SPIKE* predicted service spikes 30 minutes into future with recalls and precision of 75% and above. Also, *SPIKE* performed relatively better than other widely-used learning methods (neural nets, random forests, logistic regression).

CCS CONCEPTS

• **Software and its engineering** → *Cloud computing*; • **Computing methodologies** → Supervised learning by regression.

KEYWORDS

Cloud, optimization, data mining, parameter tuning

ACM Reference Format:

Jianfeng Chen, Joymallya Chakraborty, Philip Clark, Kevin Haverlock, Snehit Cherian, and Tim Menzies. 2019. Predicting Breakdowns in Cloud Services (with SPIKE). In *Proceedings of the 27th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '19)*, August 26–30, 2019, Tallinn, Estonia. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3338906.3340450>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ESEC/FSE '19, August 26–30, 2019, Tallinn, Estonia

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5572-8/19/08...\$15.00

<https://doi.org/10.1145/3338906.3340450>

1 INTRODUCTION

Managing cloud services is an important problem. Mismanaging such services results in server downtime and an associated loss of revenue, particularly for organizations with penalty clauses in their service contracts. Even down times of just a few hours each month can be detrimental to the professional reputation of a cloud-service provider. This is a concern since organizations with a poor reputation for reliability have a harder time attracting and retaining clients.

This paper explores one kind of breakdown—specifically, service spikes that can freeze up a cloud server. In Figure 1, such spikes are clearly visible (see the higher values). At first glance, such spikes seem relatively infrequent and very small (the y-axis of Figure 1 is in milliseconds). However, it should be remembered that a modern web page shows results from dozens of microservices, each of which uses dozens of other queries to the underlying databases. Spikes like those shown in Figure 1 can lead to frustratingly slow systems performance (e.g. very slow displays of new web pages). Hence, such spikes are critical business events that can damage an organization’s reputation as a reliable cloud service provider.

Predicting service spikes is hard since they can occur rarely and may occur as sudden extreme outliers. For example, in Figure 1, the large 11am and 8:30pm spikes might be anticipated by the steady build-up in the proceeding hour. However, between 6pm and 3am,

Figure 1: Web service response time; one cloud compute node, 11/26/2018.

we count six other spikes that are not preceded by any apparent build-up.

Another factor complicating spike prediction is the rapidly changing nature of cloud environments. For example, consider LexisNexis (the organization that funded this research). In the year or two, LexisNexis has retired its locally managed CPU farms in favor of CPU farms managed by multiple major cloud vendors. But cloud instance managements tools are evolving rapidly. Hence, like every other user of cloud-based services, LexisNexis anticipates that, in the near future, it will change its web architecture yet again.

Due to this relentless pace of change, much of the prior operational history is not relevant to current or future operations. Hence, methods that work in prior studies may not work for future studies. This means that data science teams working on cloud clusters are forced to constantly update their model.

Further, at each update, new learning technologies may be necessary. For example, the authors of this paper started with established methods for predicting service spikes [33] and when those did not work, we moved on to other methods (described later in this paper). In all, we spent three months building and discarding a dozen different predictors¹ before finding one that could handle the specifics of the LexisNexis environment.

Initially, we imagined that we would be building a recommender system that would suggest the number and type of cloud server instances that should be added or deleted in order to maintain service availability (at minimum cost). In theory, such a recommender system could be learned from the historical logs of prior nominal and off-nominal behavior.

However, once we realized how fast the cloud services were changing, we also realized that much of the historical log was no longer relevant to current practice. So we changed track and asked “what are the major pain points of running the LexisNexis cloud service?”. This new question prompted our subject matter experts to recounted various war stories about what happens when a service spike occurs. One issue with those events is gathering together the response team. “It can take five to ten minutes to realize we have a problem”, we were told, “after which it can take another few minutes of calling/texting to get everyone we need into a conference call”.

From remarks such as this, we realized our goal needed to be “early warning”. Accordingly, the goals of this project were set to be:

Build comprehensible and effective predictors for service spikes, 30 minutes into the future.

As to what constitutes a “service spike”, using the results of Section 4, we defined that to be values over 470 ms/query. Such a predictor would allow an organization to reduce their response time to forthcoming incidents. Further, in some cases, it would be possible to remove the cause the spike, thus preventing that incident from occurring in the first place.

Note that the above goal includes *comprehensible models*. Our experts required some report of the lessons learners that they can read, understand and audit. Hence, we need to use data mining methods that produce human readable models (e.g. not Naive Bayes

classifiers, not neural networks, not instance-based learners, not random forests).

1.1 Organization of this Paper

The rest of this paper explores methods for building comprehensible and effective predictors for service spikes of over 470 milliseconds per query, 30 minutes into the future. Our paper is structured as follows. Section 2 first briefly introduces some background of the LexisNexis information retrieval system, features of search engine we used, as well as the motivations of the project. Section 3 goes over all machine learning techniques explored during the project. Section 4 presents the data we collected for the prediction in the project. Section 5 reveals the exploration of all steps, leading to the managerial or technical summary. Finally we conclude the project in section 6 as well as the future work in section 7.

2 BACKGROUND

2.1 Business Context

LexisNexis provides computer-assisted legal research, business research and risk management services [1, 35]. During the 1970s, LexisNexis pioneered electronic accessibility of legal and journalistic documents [26]. Since 2006, it has the world’s largest electronic database for legal and public-records information [26].

LexisNexis provides regulatory, legal, business information and analytics to the legal community. Legal and research professionals use the LexisNexis Advance platform to find relevant information [2] that they use to prepare legal cases and drive better legal outcomes. LexisNexis operates across multiple major cloud vendors. As a large distributed system, LexisNexis is mindful of costs and wants to scale the database and search service when demand is anticipated or predicted to be needed. Hence, this paper.

2.2 System Architecture

The LexisNexis database contains over 100 billion documents and records. Records are added at the rate of nearly two million new items daily from over 50,000 sources. In all, over 20 million legal documents are processed daily from legal jurisdictions throughout the world. In addition, the databases contain over 320 million company profiles with content archives dating back 40 years.

To support all this, the Lexis Advanced product suite is a modern web application consisting of a central monolithic application, supported by hundreds of micro-services. Individual micro-services perform a wide array of functions, including content loading and enrichment, document search and retrieval, user authentication, and notifications. Many of these services are shared by multiple other services, forming an interconnected web of dependencies. When service disruptions occur, they are felt as a cascade of failures in which one service effectively disables multiple other services.

LexisNexis utilizes a number of technologies to constantly monitor the state of the application environment. Application logs, server logs, host metrics, web traffic, infrastructure status, and user activity are all monitored using various automated tools and in-house software. One of them² is primarily used to ingest, parse, and visualize data from log files, while the other tool is used to

¹During this trial-and-error period, we often took solace from Thomas Edison’s famous quote “I have not failed. I’ve just found 10,000 ways that won’t work.”

²In this paper, some terms are anonymized for proprietary business reasons.

monitor application performance metrics such as response time and throughput. Both tools allow users and automated scripts to monitor and react to quickly changing conditions in near real time.

2.3 Document Storage and Search

Currently, LexisNexis makes extensive use of some multi-model NoSQL database for document storage and searching. Such document search engine stores and queries data in documents, graph data, or relational data, providing incredible flexibility. It includes a search engine that is especially suitable for full-text searches (most documents in LexisNexis are in full-text).

The database in LexisNexis supports massive horizontal and can support installations with hundreds of nodes, petabytes of data, and billions of documents (which still processing tens of thousands of transactions per second). The system also is a *trusted database* having all the enterprise features required to handle sensitive enterprise data:

- **Advanced Security:** It offers granular security controls at the document and even element/property level, redaction and anonymization for safe sharing, as well as advanced encryption.
- **ACID Transactions:** It has multi-document ACID transactions that provide data consistency even with large-scale transaction applications.
- **Cloud Neutral:** It has been successful running in the cloud for over a decade and it is compatible with any public cloud provider.

However, such trusted database can lead to operational issues. At its core, it is a *twin instance* system where data is stored on Node1 and indexed on Node2. This twin architecture increases survivability of the system against insult (if the index node goes down, it can be rebuilt elsewhere). On the other hand, experience has shown that this twin model can complicate the operation of such database system. Nodes cannot be simply added (if more performance is required) or removed (if we want to save operational costs if the CPUs are under-utilized). There is also an additional operational cost of such a system—after a crash (when data has to be rebuilt), some expert human supervision is required to appropriately partition the data across the servers.

In the meanwhile, LexisNexis may soon be running its database servers on some open-sourced container orchestration – at which time, the expertise needed to run a database server will need to be updated, yet again. Because of this relentless pace of change in cloud services, we made the design decision to build *SPIKE* without using detailed knowledge of the current internal database system. Instead, as discussed in section 4, when we did data collection, we restricted ourselves to measurements we might reasonably expect to see in a wide range of future cloud environments.

3 DATA MINING TECHNOLOGY

We adopted four widely used machine learning models as system spike learners, including logistic regression (LG), classification and regression trees (CART)/random forest (RF), artificial neural network (ANN) and long-short term memory network (LSTM). We chose these four learners since a recent survey of predicting service spike [33] listed LG, CART and ANN as the most common machine

learning models to predict the web server response time. They also suggest several other ML approaches applicable to time series modeling, including some more complicated ANNs developed and used in deep learning (e.g., LSTM).

3.1 Logistic Regression

Logistic regression analyzes the relationship between multiple independent variables and a categorical dependent variable, and estimates the probability of occurrence of an event by fitting data to a logistic curve [28]. Two kinds of logistic regression are binary or binomial logistic regression and multinomial logistic regression:

- *Binary logistic regression* is used when the dependent variable is dichotomous and the independent variables are either continuous or categorical.
- When the dependent variables are not dichotomous and is comprised of more than two categories, a *multinomial logistic regression* can be employed.

Unlike linear regression, logistic regression can directly predict probabilities, i.e. the odds of a dependent variable happens. The most essential assumption to the logistic regression is that the independent variables are linearly related to the log odds, that is the logit of the probability defined as $\text{logit}(p) = \log \frac{p}{1-p}$ ($0 < p < 1$).

Logistic regression is used in various fields, including the system maintenance. For example, Hoffert *et al.* [21] trained the logistic regression models to predict the response time of a search and rescue (SAR) operations system and therefore simplify the configuration of middleware and adaptive transport protocols.

3.2 CART

The decision tree, or specifically, the classification and regression trees *CART* [30] is another common type of supervised learning algorithm. It works for both categorical and continuous input and output variables. CART splits the population or sample into two or more homogeneous sets (or sub-populations) based on most significant splitter/differentiator in input variables. The leaf nodes of the tree contain an dependent variable which is used to make a prediction.

There are two types of trees:

- *Classification Tree* which serves problems with categorical target variables;
- *Regression Tree* which is applicable to problems with continuous target variables.

In our work, when applying the decision tree, we treated the target (service spikes indicator) as continuous target, instead of binary category. An estimation of service spike values let operation team staffs have the sense how urgent the further actions should take.

In terms of generating comprehensible models, classification and regression trees are our preferred choice. In our experience, if they can be kept under a few dozen nodes, decision/regression trees are fast to read and understand. Using the trees, engineers can figure out the key factors that contribute the targets. That information is constructive in further system refactoring. For example, we can adjust more storage and bandwidth resource if the I/O is the key factor to system performance to some microservice.

3.3 Other Learners

In order to assess the impact of our “comprehensability” requirement on learner performance, as described in this section, we also explored some learning models that can produce somewhat opaque results.

Random Forests construct multiple trees at the training time. The prediction of such a forest comes from the majority view of all its trees. Random forest was first introduced by Tin Kam Ho [19] and has been applied in many ML applications [4, 5, 17]. Random forests can produce a very large set of trees which can be hard to read and understand.

Neural networks consist of input and output layers, as well as (in most cases) a hidden layer consisting of units that transform the input into something that the output layer can use [11]. The most basic type of neural net is something called a *feedforward neural network*, in which information travels in only one direction from input to output. A more widely used type of network is the *recurrent neural network* (RNN), in which data can flow in multiple directions.

Neural network models represent their knowledge in a somewhat arcane distributed manner which is not human comprehensible.

Long-short term memory networks [20], or LSTM for short, is a special kind of RNN, capable of learning long-term dependencies. Figure 2 shows a simple recurrent network. In RNN, the model is interpreted not as cyclic, but rather as a deep network with one layer per time step and shared weights across time steps. This algorithm is called back propagation through time [39]. Recurrent Neural Networks suffer from problems with short-term memory. During back propagation, recurrent neural networks suffer from the vanishing gradient problem (gradients are values used to update a neural networks weights). The vanishing gradient problem is when the gradient shrinks as it back propagates through time. If a gradient value becomes extremely small, it does not contribute too much learning. LSTM was created as the solution to this short-term memory problem. The long short-term memory block is a complex unit with various components such as weighted inputs, activation functions, inputs from previous blocks and eventual outputs. The block is called a long short-term memory block because RNN is using a structure founded on short-term memory processes to create longer-term memory.

3.4 Data Pre-Processing with SMOTE

As mentioned above, our training data is very imbalanced. Specifically, from Nov 2018 to Jan 2019, the service spikes happened at only 3.4% of the time. Such imbalanced training data makes it hard for classification models to detect rare events [32].

There are several ways to apply *resampling* to mitigate for class imbalance [10, 25, 37, 38]:

- Over-sampling to make more of the minority class;
- Under-sampling to remove majority class items;
- Some hybrid of the first two.

Machine learning researchers [18] advise that under-sampling can work better than over-sampling if there are hundreds of minority observations in the datasets. When there are only a few dozen minority instances, over-sampling approaches are superior

Figure 2: A Simple Recurrent Neural Network [24]

to under-sampling. In the case of large size of training samples, the hybrid methods would be a better choice.

The Synthetic Minority Oversampling Technique (SMOTE) [10] is a hybrid algorithm that performs both over- and under-sampling. SMOTE calculates the k nearest neighbors for each minority class samples. Depending on the amount of oversampling required, one or more of the k -nearest neighbors are picked to create the synthetic samples. This amount is usually denoted by oversampling percentage (e.g., 50% by default). The next step is to randomly creating a synthetic sample along the line connecting two minority samples.

3.5 Parameter Tuning with Differential Evolution

In machine learning, model hyperparameter are values in machine learning models that can require different constraints, weights or learning rates to generate different data patterns, e.g., the number of neighbours in k -Nearest Neighbours (KNN) [22]. Such hyperparameters are very important because they directly control the behaviors of the training algorithm and also impact the performance of the model being trained. Therefore, choosing appropriate hyperparameters plays a critical role in the performance of machine learning models. Hyperparameter tuning is the process of searching the most optimal hyperparameters for machine learning learners [8, 12].

Recent studies have shown that hyperparameter optimization can achieve better performance than using “off-the-shelf” configurations in several research areas in software engineering, e.g., software defect prediction [14, 15, 23, 27, 34] and software effort estimation [40]. To the best of our knowledge, we are first to apply hyperparameter optimization in response time prediction.

Hyperparameter optimization can be implemented in many ways:

- *Grid search* [7] loops through all combinations of all parameters. Although Grid search is a simple to implement, it suffers if data have high dimensional space called the “curse of dimensionality”. Previous work has shown that grid

Table 1: List of hyperparameters tuned.

Learner	Parameters	Default	Best	Descriptions
CART	min samples split	2	0.61	The minimum number of samples required to split an internal node.
	max depth	None	5	The maximum depth of the tree.
RF	n_estimators	10	20	The number of trees in the forest.
	min samples split	2	0.1	The minimum number of samples required to split an internal node.
	max depth	None	3	The maximum depth of the tree.

search might miss important optimization [16] and is a time-wasting process since only a few of the tuning parameters really matters [6].

- *Random search* [6] randomly samples the search space and evaluates sets from a specified probability distribution. Such random searches do not use information from prior experiments to select the next set and also it is very difficult to predict the next set of experiments.
- *Bayesian optimization* [29] works by assuming the unknown function was sampled from a Gaussian Process and maintains a posterior distribution for this function as observation are made. However, it might be best-suited for optimization over continuous domains with small number of dimensions [13].

This paper uses *Differential evolution* (DE) for hyperparameter optimization. DE has proven useful in prior SE tuning studies [15]. Also, our reading of the current literature is that there are many advocates for differential evolution like Vesterstrom *et al.* [36] showed DE to be competitive with particle swarm optimization and other genetic algorithms.

The premise of DE is that the best way to mutate the existing tunings is to extrapolate between current solutions. DE builds a population P from a small number of randomly selected solutions of size np . Then, each member of the population is compared against a *mutant* built as follows. Three solutions $(a, b, c) \in P$ are selected at random. For each tuning parameter k , at some probability cr , we replace the old tuning x_k with y_k as $y_k = a_k + f \times (b_k - c_k)$ where f is a parameter controlling differential weight. The main loop of DE runs over the population P of size np , replacing old items with new candidates (if new candidate is better). This means that, as the loop progresses, the population is full of increasingly more valuable solutions (which, in turn, helps extrapolation).

For pragmatic reasons we did not tune all parameters of all learners. LSTMs took 30 minutes to test each tuning. Given our DE settings, that would have required 6000 hours of CPU; i.e. 25 weeks. Table 1 shows the parameters that we did tune. During that tuning process, we asked our optimizers to maximize recall and precision. As to the control parameters of DE, using advice from Storn and Fu *et al.* [14, 31], we set $\{np, f, cr\} = \{20, 0.75, 0.3\}$. Also, the number of generations gen was set to 10 to test the effects of a very CPU-light optimizer.

For the set of parameters we did tune, see Table 1.

4 DATA

For this analysis, we collected data from the LexisNexis N-document database searching microservice. This N-document database contains 20+ million documents.

Table 2: Statistics of monitored data after the final steps in the pipeline.

Percentile	EE	AS	MP	THR
0% (min)	0	0.53	67800	2.83
25%	0.12	0.96	105000	15.90
50%	1.72	0.97	107000	35.60
75%	3.88	0.97	111000	102
100% (max)	347	0.99	272000	243

Recall from the introduction that the technology used in LexisNexis' cloud systems is changing rapidly. Accordingly, when we did data collection at one-minute intervals, and restricted ourselves to measurements we might reasonably expect to see in a wide range of future cloud environments.

We collected data on:

- the current *web service response time* of the N-document service;
- that same response time at $t \in \{5, 10, 15, 30, 60, 90, 120, 150, 180, 300, 1440\}$ minutes into the past.

We also collected the following data on the for each 13 microservices directly upstream on the N-document services

- (1) *Response time* for service x ;
- (2) *All Logged Errors Per Minute* (EE_t);
- (3) *Total Physical Memory Used in MB* (MP_t);
- (4) *Web Transaction Throughput* (THR_t).
- (5) The *application performance index Apdex Score* (AS_t). The Apdex score is defined as the number of satisfied samples plus half of the tolerating samples plus none of the frustrated samples, divided by all the samples in unit minute. That is:

$$Apdex = (Satisfactory\ samples + 0.5 \times Tolerating\ samples + 0 \times Frustrated\ samples) / Total\ samples$$

Here, *satisfactory*, *tolerating* and *frustrating* are defined in the standard way, as per [3].

The columns of our collected data had the distributions of Table 2. Figure 3 and 4 show the histograms of attributes (independent features) and class (dependent feature) respectively. From these Figure 3 and 4, we have the following observations:

- All our variables are highly non-evenly distributed, i.e. all of them have large standard deviation.
- In statistics, a *long tail* of some distributions of numbers is the portion of the distribution having a large number of occurrences far from the “head” or central part of the distribution [9]. We see this pattern among all attributes.
- From Figure 4 we can see that in majority time, the system response time varied from 0-470. By selecting for $x > 470$,

Figure 3: Independent features histogram. x -axis shows sorted all monitoring values. y -axis is the corresponding frequency.

Figure 4: Histogram of Web Service Response Time. x -axis is the service response time in milliseconds. y -axis is the frequency in the training set. Note that this plot shows a *dependent* variable that is *different* to the THR_t independent variable discussed in the text.

we could focus this study on the most outstanding service spikes.

5 PREDICTING WITH SPIKE

LexisNexis is serving customers from hundreds of areas. Their behavior pattern may be changing at all times, therefore lead to various situations of service spikes. For example, the way customers from Wall Street using the LexisNexis service is different to those from Silicon Valley – at least, they are interested in distinct group of documents. As a result, model to predict "financial news" spikes is not applicable to predict "tech scandal" spikes. Therefore it is important to train *SPIKE* on recent local data that is specific to a particular web service.

That said, deciding what data mining method to apply is a time-consuming and CPU expensive process. Therefore, our work was divided into two stages:

- (1) *Model selection and tuning* where we pruned options to find one promising model. For this stage, we used one month of data divided into an 80% train and 20% test phase (where the test data was selected from the last week of the month).
- (2) *Testing our most promising method*. For this stage, we used a different month of data to test the method selected during stage one. We stepped through this data in "windows" of ten minutes. At each step i , we trained the model (found by *model selection and tuning*) using the next 24 hours of data (i.e. windows i to $i + 144$). This model was then tested using data from the next half hour of data (i.e. windows $i + j, j \in \{145, 146, 147\}$). This means that the stage2 results (reported below) come from 4317 different train/test pairs.

For details on these two stages, see below.

5.1 Stage One: Model Selection and Tuning

In this stage, we sorted our one month data by time then trained on the first 80% and tested on the last 20%³. The models were trained to get the real-time *web service response time* in this stage.

Not all treatments were applied to all data sets. For example, as mentioned above, the neural nets were too slow to tune. As for the other learners, when optimizing hyperparameters, at the request of our business users, we optimized for maximizing recall. Recall is defined for a two-class classifier so if this stage, to guide the hyperparameter optimization, we defined "spike" as per Figure 4; i.e. greater than 470ms.

Important point: SMOTE or hyperparameter optimization never used information from the test data. If we applied SMOTE or hyperparameter optimization, these algorithms were used *only* on the training data; i.e. our results are not over-fitted to test data.

In all, we explored the 14 treatments of Table 3. As shown by the rows highlighted in gray, three methods achieved very high

³We did not use a randomized strategy to produce train/test sets since it makes more methodological sense to predict the future from the past.

Table 3: Learning on first 80% of data, testing on the more recent 20%. Results sorted by recall (and *higher values are better*). ANN=Artificial Neural network, LSTM=Long-short Term Memory. In this table *higher values for recall, precision*.

			TP	FP	FN	TN	$r = \frac{TP}{TP+FN}$	$p = \frac{TP}{TP+FP}$
Learner	SMOTE?	TUNE?	TP	FP	FN	TN	Recall(%)	Precision(%)
ANN(10×4)			504	1756	27	799	95	22
Regression Tree (CART)	✓	✓	493	868	38	1687	93	36
Random Forest	✓		489	764	42	1791	92	39
Random Forest			428	439	103	2116	81	49
Random Forest	✓	✓	417	437	114	2118	79	49
Decision Tree (CART)	✓		294	432	237	2123	55	40
Decision Tree (CART)			232	362	299	2193	44	39
LSTM	✓		53	78	478	2477	10	40
Logistic Regression	✓		42	36	489	2519	8	54
ANN(5×4)	✓		40	44	491	2511	8	48
ANN(5×4)			35	12	496	2543	7	74
Logistic Regression			31	9	500	2546	6	77
LSTM			30	23	501	2532	6	57
ANN(10×4)	✓		31	33	500	2522	6	48

recalls of over 90% (Random Forests, CART, ANN). Initially, off-the-shelf CART performed poorly. However, when augmented with SMOTE and tuning, CART achieved very high recalls of over 90% (the associated precisions are not good– which is a problem solved by the sensitivity analysis of the next section).

Applying the *comprehensibility* criteria, our summary of Table 3 is that SMOTE+ Regression Trees (CART) + hyperparameter optimization performs best. While ANN offered marginally better recalls, it is hard to read those models. CART, on the other hand, produced the simple regression tree of Figure 5.

Since this tree is easy to comprehend, it is easy to extract important business knowledge about LexisNexis cloud machines. For example:

- Among all 13 nodes studied, only three were found to be important by this tree: ret, lo and sh-synr. Prior to this

study, the importance of these nodes to healthy operations at LexisNexis has not been realized.

- To avoid spikes, engineers are advised to take action that avoids the the red branches of Figure 5.

5.2 Stage Two: Testing our Most Promising Method

Stage one found that the tree learner (CART) with specific hyperparameters (as shown in Table 1) was comparatively better than several other methods.

This second stage tests if that model is useful on real-world data. To predict the *web service response time* at the moment ($t+.5$)hr, SPIKE trained the data from $[t - L\text{hrs}, t)$, where L is the training window size. We explored $L \in \{0.25, 0.5, 1, 5, 24, 48\}$ and found that best results come from training from the last $L = 24$ hrs of prior data.

Figure 5: Model generated from CART ret, lo and sh-synr are node names. Branches leading to spikes (with service response time >470ms/query) are highlighted in red.

Figure 6: Precision and recalls under different sensitivity to predict services spikes half an hour ahead. x -axis is the threshold value used to trigger an alarm.

We found that the precisions and recalls achieved by CART+SMOTE+ optimization were sensitive to our threshold for recognizing a spike. Hence, we show results where we declared a “spike” being defined as more than from 370ms to 490ms.

Figure 6 shows results seen while adjusting the threshold for predicting a spike. As shown in this figure, at a threshold of 404ms, can achieve precisions and recalls of 75% or higher.

It is worth mentioning that, according to the experiment, there always exist some false positives or false negatives. The impact of false positive and false negative signals from the on the business is generally limited, but it can grow to devalue to the tool. High levels of false negative signals would result in no action taken by support engineers prior to service incidents. This would essentially be “business as usual”, and likely not a negative impact other than decreasing confidence in the tool. On the other hand, high levels of false positive signals would cause support engineers to engage in monitoring activities prematurely, often diverting them from other tasks. While this seems like a high cost, discussions with LexisNexis site reliability engineering team indicates that this cost is acceptable. The time it takes to open an application performance monitoring tool to view current status is a negligible cost in comparison to the benefit of potentially avoiding a customer impacting service degradation.

6 CONCLUSIONS AND LESSONS LEARNED

Using regression trees (e.g. CART) and synthetic over-sampling of rare events (e.g. via SMOTE) and hyperparameter optimization (e.g. using DE), optimizers it is possible to build effective and comprehensive predictors for service spikes. *SPIKE* can predict with reasonable recall and precisions if a spike will occur in the next 30 minutes. Further, *SPIKE* can report its reasoning via a very small and easily comprehensible tree, from which we can learn important (and previously unknown) aspects about this domain.

The factors that lead to service spikes are highly context specific. Much time was spent in this work trying solutions from other sites [33], which proved less-than-satisfactory for this problem (neural nets, logistic regression). There exist tools for exploring a large number of options within data miners. If we had our time again, we would *first* commission a hyperparameter optimization (a tool

To understand the precision and recalls, consider the threshold 370ms that triggers an alert:

- Among all actual future spikes (response time > 470ms), 49.2% of them have predicted value > 370ms (therefore trigger the alarm), which shows as point A;
- Among all alarms triggered (i.e. predicted future response time > 370ms), 93.8% of them did exceed the threshold (having response time > 470ms), indicated as point B.

to explore all those options) before *secondly* use those optimizers to faster explore different data mining options.

But even with hyperparameter optimizers, building predictors is a complex tasks (certainly, much more than running one query, then glancing at a simple data dashboard). There is considerable creativity required in how to design the inputs to a learning problem and how to find tune the resulting models. For example, in this work, we made poor progress until we somewhat serendipitously decided to:

- Add inter-node information to the training set (see Section 4);
- Conduct a sensitivity analysis (see previous section).

More generally, a modern cloud environment can generate petabytes of operational logs, every day. For example, LexisNexis constantly monitors the state of its cloud services, collecting data from many microservices at one-minute intervals. A data science team exploring the problem of service spikes needs considerable business knowledge to “slice and dice” the data. In all, the results of this paper took three months to generate:

- 1 month of a LexisNexis data engineer generating our training data by writing complex joins across large datasets.
- 1 month of inductive engineering, applying different data mining methods to the data. As mentioned above, this proved to be a tedious task that required developed and discarding a dozen very bad predictors before finding one that achieved useful results
- 1 month of a senior LexisNexis engineer serving as a liaison between our team and the rest of LexisNexis. The importance of the liaison cannot be overstated. That person (a) maintained senior management’s awareness and enthusiasm for this project; (b) organized access to numerous subject matter experts.

When staffing similar efforts in the future, we recommend a similar “three-sided” team comprising inductive engineers, data engineers, and business knowledge experts.

7 FUTURE WORK

We believe *SPIKE* is a general method for managing rare, but critical, CPU issues in complex cloud environments:

- (1) For each node, *SPIKE* trains models using (a) *intra-node* details about the recent history of that node as well as (b) some *inter-node* knowledge about connected nodes.
- (2) For rare events, it is important to use class rebalancing tools (like SMOTE).
- (3) Also, since the factors that lead to service spikes are highly context specific, it is useful to employ hyperparameter optimization (like DE).

In theory, these three principles should apply to other services at LexisNexis and other organizations. In future work we aim to test that conjecture using more data.

Furthermore, most tools SRE's deal with either have an "Alarm" state or an "Ok" state. This type of tool is more a "probability" tool of a potential alarm state (defined as the threshold like Figure 6). It is an interesting future work to call out a "confidence" that something is going to go into an "Alarm" state.

Also, after *prediction* comes *diagnosis* and *repair*. If we build trees like Figure 5 from more data (covering more months and more LexisNexis services) then we would be able to uncover critical thresholds for critical nodes that most effect LexisNexis services. Using that knowledge, plus more subject matter expertise, we should then be able to propose spike reduction policies.

ACKNOWLEDGEMENTS

This work was partially funded by (a) a gift from LexisNexis managed by Phillpe Poignant; and (b) an NSF CCF grant #1703487.

REFERENCES

- [1] 1994. Company News; A Name Change is Planned for Mead Data Central. *The New York Times*. (1994). <https://www.nytimes.com/1994/12/02/business/company-news-a-name-change-is-planned-for-mead-data-central.html?src=pm>
- [2] 2019. LEGAL. *www.relx.com* (2019). <https://www.relx.com/our-business/market-segments/legal>
- [3] 2019 (accessed Apr 1, 2019). *Apdex Alliance*. <http://www.apdex.org/>.
- [4] Daniel C Alexander, Darko Zikic, Jiaying Zhang, Hui Zhang, and Antonio Criminisi. 2014. Image quality transfer via random forest regression: applications in diffusion MRI. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 225–232.
- [5] Mariana Belgio and Lucian Drăguț. 2016. Random forest in remote sensing: A review of applications and future directions. *ISPRS Journal of Photogrammetry and Remote Sensing* 114 (2016), 24–31.
- [6] James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *Journal of Machine Learning Research* 13, Feb (2012), 281–305.
- [7] James S Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. 2011. Algorithms for hyper-parameter optimization. In *Advances in neural information processing systems*. 2546–2554.
- [8] Andre Biedenkapp, Katharina Eggensperger, Thomas Elsken, Stefan Falkner, Matthias Feurer, Matilde Gargiani, Frank Hutter, Aaron Klein, Marius Lindauer, Ilya Loshchilov, et al. 2018. Hyperparameter Optimization. *Artificial Intelligence* 1 (2018), 35.
- [9] Alpheus Bingham and Dwayne Spradlin. 2011. The Long Tail of Expertise. *Pearson Education* ISBN 9780132823135 (2011). <https://books.google.com/books?id=9qLfsonmwhAC&pg=PT5&v=onepage&q&f=false>
- [10] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16 (2002), 321–357.
- [11] Luke Dormehl. 2019. What is an artificial neural network? Here's everything you need to know. *EMERGING TECH* (2019). <https://www.digitaltrends.com/cool-tech/what-is-an-artificial-neural-network/>
- [12] Luca Franceschi, Michele Donini, Paolo Frasconi, and Massimiliano Pontil. 2017. Forward and reverse gradient-based hyperparameter optimization. *arXiv preprint arXiv:1703.01785* (2017).
- [13] Peter I Frazier. 2018. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811* (2018).
- [14] Wei Fu and Tim Menzies. 2017. Easy over hard: A case study on deep learning. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*. ACM, 49–60.
- [15] Wei Fu, Tim Menzies, and Xipeng Shen. 2016. Tuning for software analytics: Is it really necessary? *Information and Software Technology* 76 (2016), 135–146.
- [16] Wei Fu, Vivek Nair, and Tim Menzies. 2016. Why is Differential Evolution Better than Grid Search for Tuning Defect Predictors? *arXiv preprint arXiv:1609.02613* (2016).
- [17] Ulrike Grömping. 2009. Variable importance assessment in regression: linear regression versus random forest. *The American Statistician* 63, 4 (2009), 308–319.
- [18] Guo Haixiang, Li Yijing, Jennifer Shang, Gu Mingyun, Huang Yuan Yue, and Gong Bing. 2017. Learning from class-imbalanced data: Review of methods and applications. *Expert Systems with Applications* 73 (2017), 220–239.
- [19] Tin Kam Ho. 1995. Random decision forests. In *Proceedings of 3rd international conference on document analysis and recognition*, Vol. 1. IEEE, 278–282.
- [20] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [21] Joe Hoeffert, Daniel Mack, and Douglas Schmidt. 2009. Using machine learning to maintain pub/sub system qos in dynamic environments. In *Proceedings of the 8th international workshop on adaptive and reflective middleware*. ACM, 4.
- [22] James M Keller, Michael R Gray, and James A Givens. 1985. A fuzzy k-nearest neighbor algorithm. *IEEE transactions on systems, man, and cybernetics* 4 (1985), 580–585.
- [23] Rahul Krishna, Tim Menzies, and Lucas Layman. 2017. Less is more: Minimizing code reorganization using XTREE. *Information and Software Technology* 88 (2017), 53–66.
- [24] Zachary C. Lipton, John Berkowitz, and Charles Elkan. 2015. A Critical Review of Recurrent Neural Networks for Sequence Learning. *arXiv:cs.LG/1506.00019*
- [25] Inderjeet Mani and I Zhang. 2003. kNN approach to unbalanced data distributions: a case study involving information extraction. In *Proceedings of workshop on learning from imbalanced datasets*, Vol. 126.
- [26] Stephen Miller. 2012. For Future Reference, a Pioneer in Online Reading. *The Wall Street Journal* (2012). <https://www.wsj.com/articles/SB10001424052970203721704577157211501855648?KEYWORDS=lexisnexis>
- [27] Haidar Osman, Mohammad Ghafari, and Oscar Nierstrasz. 2017. Hyperparameter optimization to improve bug prediction accuracy. In *Machine Learning Techniques for Software Quality Evaluation (MaLTesQuE), IEEE Workshop on*. IEEE, 33–38.
- [28] Hyeoun-Ae Park. 2013. An Introduction to Logistic Regression: From Basic Concepts to Interpretation with Particular Attention to Nursing Domain. *Korean Society of Nursing Science* 43, 2 (2013), 1–5. <http://synapse.koreamed.org/DOIx.php?id=10.4040%2Fjkan.2013.43.2.154>
- [29] Martin Pelikan, David E Goldberg, and Erick Cantú-Paz. 1999. BOA: The Bayesian optimization algorithm. In *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation-Volume 1*. Morgan Kaufmann Publishers Inc., 525–532.
- [30] Leszek Rutkowski, Maciej Jaworski, Lena Pietruczuk, and Piotr Duda. 2014. The CART decision tree for mining data streams. *Information Sciences* 266 (2014), 1–15.
- [31] Rainer Storn and Kenneth Price. 1997. Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization* 11, 4 (1997), 341–359.
- [32] Yanmin Sun, Andrew KC Wong, and Mohamed S Kamel. 2009. Classification of imbalanced data: A review. *International Journal of Pattern Recognition and Artificial Intelligence* 23, 04 (2009), 687–719.
- [33] Yang Syu, Chien-Min Wang, and Yong-Yi Fanjiang. 2018. A Survey of Time-Aware Dynamic QoS Forecasting Research, Its Future Challenges and Research Directions. In *International Conference on Services Computing*. Springer, 36–50.
- [34] Ayse Tosun and Ayse Bener. 2009. Reducing false alarms in software defect prediction by decision threshold optimization. In *Proceedings of the 2009 3rd International Symposium on Empirical Software Engineering and Measurement*. IEEE Computer Society, 477–480.
- [35] Ashlee Vance. 2010. Legal Sites Plan Revamps as Rivals Undercut Price. *The New York Times*. (2010). https://www.nytimes.com/2010/01/25/technology/25westlaw.html?_r=1&ref=reedelsevier
- [36] Jakob Vesterstrøm and Rene Thomsen. 2004. A comparative study of differential evolution, particle swarm optimization, and evolutionary algorithms on numerical benchmark problems. In *Evolutionary Computation, 2004. CEC2004. Congress on*, Vol. 2. IEEE, 1980–1987.
- [37] James Walden, Jeff Stuckman, and Riccardo Scandariato. 2014. Predicting vulnerable components: Software metrics vs text mining. In *Software Reliability Engineering (ISSRE), 2014 IEEE 25th International Symposium on*. IEEE, 23–33.
- [38] Byron C Wallace, Thomas A Trikalinos, Joseph Lau, Carla Brodley, and Christopher H Schmid. 2010. Semi-automated screening of biomedical citations for systematic reviews. *BMC bioinformatics* 11, 1 (2010), 55.
- [39] P. J. Werbos. 1990. Backpropagation through time: what it does and how to do it. *Proc. IEEE* 78, 10 (Oct 1990), 1550–1560. <https://doi.org/10.1109/5.58337>
- [40] Tianpei Xia, Rahul Krishna, Jianfeng Chen, George Mathew, Xipeng Shen, and Tim Menzies. 2018. Hyperparameter Optimization for Effort Estimation. *arXiv preprint arXiv:1805.00336* (2018).