

## A2Log: Attentive Augmented Log Anomaly Detection

Thorsten Wittkopp<sup>1</sup>, Alexander Acker<sup>1</sup>, Sasho Nedelkoski<sup>1</sup>, Jasmin Bogatinovski<sup>1</sup>,  
Dominik Scheinert<sup>1</sup>, Wu Fan<sup>2</sup> and Odej Kao<sup>1</sup>

<sup>1</sup>Technische Universität Berlin, Germany, {t.wittkopp, nedelkoski, firstname.lastname}@tu-berlin.de

<sup>2</sup>Huawei, China, {wufan3@huawei.com}

### Abstract

*Anomaly detection becomes increasingly important for the dependability and serviceability of IT services. As log lines record events during the execution of IT services, they are a primary source for diagnostics. Thereby, unsupervised methods provide a significant benefit since not all anomalies can be known at training time. Existing unsupervised methods need anomaly examples to obtain a suitable decision boundary required for the anomaly detection task. This requirement poses practical limitations. Therefore, we develop A2Log, which is an unsupervised anomaly detection method consisting of two steps: Anomaly scoring and anomaly decision. First, we utilize a self-attention neural network to perform the scoring for each log message. Second, we set the decision boundary based on data augmentation of the available normal training data. The method is evaluated on three publicly available datasets and one industry dataset. We show that our approach outperforms existing methods. Furthermore, we utilize available anomaly examples to set optimal decision boundaries to acquire strong baselines. We show that our approach, which determines decision boundaries without utilizing anomaly examples, can reach scores of the strong baselines.*

### 1. Introduction

IT systems are rapidly evolving to meet the growing demand for new services and applications in various economic fields. Many companies outsource their services to the cloud [1]. This outsourcing entails growing data centers, increasingly large networks, and interconnected devices to provide the required IT services. Despite accelerating innovations and business opportunities, this trend increases complexity and thus, aggravates the operation and maintenance of these services and systems [2]. Operators of the services require assistance in maintaining control of

this complexity to ensure dependability, stability, and serviceability. The field of artificial intelligence for IT operations (AIOps) is intended to support service and system operators to meet these challenges. Thereby, anomaly detection is a vital part, which is applied to monitoring data such as metrics, logs, or traces.

Log data is a primary source for troubleshooting since they record events during the execution of service applications. These logging events evolve due to software updates. Therefore unsupervised methods are valuable due to their sense of new and unknown anomalies [3, 4]. The second benefit of unsupervised methods is that they do not require labeled data, which is hard to obtain and cost-intensive [5]. Due to the complexity of IT services, the log data volume is growing to the extent that it cannot be manually analyzed. Therefore, recent research uses deep learning to analyze log data and perform anomaly detection [6, 7, 8]. These studies mostly assume the existence of a sufficient amount of labeled validation data for parameter tuning. However, in production settings, where services evolve, such data is hard to obtain, volatile, and requires manual evaluation by experts. In particular, when it comes to deriving an anomaly decision, these methods lack the capability to obtain a decision in an unsupervised manner. The anomaly decision in all mentioned methods is based on a decision boundary that decides on the respective binary class. Commonly employed methods need to be aware of the anomalies in the validation data to set this decision boundary optimally. This relatively strong requirement poses limitations.

In this paper, we address these problems via a two-fold solution called *A2Log*. First, we create a neural network based on the self-attention mechanism to perform an anomaly scoring. Second, we perform data augmentation to generate deviations on the respective training data [9, 10], to analyze the model response and calculate the final decision boundary. In both steps, only normal training data is used, thus we call our approach unsupervised.

The contributions of this work contain the following:

- An unsupervised anomaly detection method for log data, based on an encoder transformer architecture.
- An unsupervised decision boundary calculation for the anomaly decision, based on a novel data augmentation method for log data.
- An evaluation of the method, based on three different real-world and publicly available datasets, which are BGL, thunderbird, and spirit<sup>1</sup>.
- An evaluation of the method, based on an industry dataset from an IT service provider.

This paper is structured as follows. First, related work is presented in section 2. A general framework for anomaly detection on log data as well as a problem description for the final anomaly decision are presented in section 3. Preliminaries and implementation details of our method are presented in section 4. Finally, section 5 presents our evaluation and section 6 is concluding the work.

## 2. Related Work

Detecting abnormal events in large-scale systems, indicated by log files, is crucial for creating dependable services. Therefore, log analysis becomes increasingly important for industry and academia [11] and a wide range of different anomaly detection techniques have been developed and discussed in detailed surveys [4, 12]. They utilize different forms of log templates and log embeddings to convert the logs to a machine-readable format. Commonly used anomaly detection methods are support vector machines [13] and principal component analysis [14]. In addition, there are rule-based [15], tree-based [16], statistical [17], as well as methods based on clustering [18, 19, 20, 21]. Whereas recent anomaly detection methods are mainly designed with neural networks [22, 23, 24] and based on encoder architectures [25, 26]. Equally, recent methods utilizing the attention mechanism that is often used in encoder architectures [27].

We classify anomaly detection methods in two different types: supervised and unsupervised [28]. Supervised methods are usually more accurate, though they train the anomalies of the specific dataset as well [22, 29]. However, not all anomalies can be known in advance. Hence, in industrial applications, unsupervised methods are more practical as anomaly labels are mostly unavailable [30, 5].

<sup>1</sup><https://www.usenix.org/cfdr-data>

Consequently, we focus on unsupervised methods in this paper. Several unsupervised learning methods, based on neural networks, have been proposed, of which we present a selection in the following.

Du et al. [7] proposed DeepLog, a Long short-term memory (LSTM) network architecture that is capable of identifying abnormal sequences of log messages. For this, log templates are generated and sequences of templates are formed as model inputs. The model provides a ranked output with probabilities for the next template in a given sequence. The anomaly detection is then based on whether the next template has a high probability or not.

LogAnomaly [30] is similar to DeepLog and predicts the next log message in a sequence of log messages. Instead of utilizing sequences of log templates, LogAnomaly utilizes sequences of log embeddings to improve prediction effectiveness.

Yang et al. [31] combining the attention mechanism with a gated recurrent network architecture to perform anomaly detection on log data. Thereby the log messages are transformed into log templates, to then predict if a sequence of log templates is normal or abnormal. Furthermore, the labels of the training data are estimated to incorporate knowledge on historical anomalies into the model.

Another approach is Logsy as proposed by Nedelkoski et al. [8]. This approach also incorporates the attention mechanism with an encoder architecture, where log embeddings are calculated. The embeddings of normal log messages are condensed into a centroid using a hyperspherical loss function, whereby embeddings of abnormal log messages are pushed away. The anomaly detection task is then based on the distance to the centroid

Guo et al. [6] introduce LogBERT. It utilizes the transformer network from BERT [32], which consists of an encoder and a decoder, including the attention mechanism. Like DeepLog, it tries to predict a targeted log template of a sequence. Therefore, they utilize temporal-related log embeddings around the prediction target as inputs. To predict the targeted log template, it utilizes Cross-Entropy-Loss, extended by a hyperspherical loss function, to ensure compactness of the embeddings.

However, all methods presented above use a manually set decision boundary for anomaly detection. However, optimization requires knowledge about the anomalies, which limits these methods.

### 3. Anomaly Detection on Log Data

In this section, we describe the general framework for anomaly detection on log data. Afterwards, we describe the problem of automatically finding an accurate decision boundary for the anomaly detection task.

#### 3.1. General Framework

Anomalies are patterns in data that do not conform to a defined notion of normal behavior [4]. Anything that deviates from normal behavior can be considered abnormal behavior. Steinwall et al. [33] state that this can be considered as a binary classification task. Consequently, anomaly detection on log data is defined as the problem of assigning a binary label to each log message. There are two principal approaches to anomaly detection: Supervised and unsupervised. In terms of anomaly detection on log data, supervised means that both classes, normal and abnormal, are used during the training phase. Unsupervised means that the model is trained on normal log messages only. As not all possible anomalies in log data can be known and used for training [34], unsupervised approaches are well suited for log data scenarios, and thus are of high interest for industry and academia [11, 3]. Therefore, the challenge is to develop a good understanding of normal log messages, e.g. by internalizing their usual structure, such that any significant deviations can be treated as log messages representing abnormal behavior. Thus, a general framework for unsupervised approaches can be described as follows:

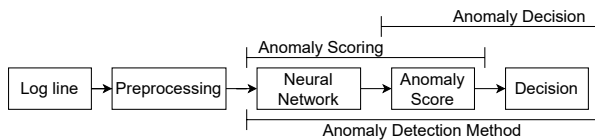


Figure 1. Structure of an unsupervised anomaly detection method for log data content.

Figure 1 illustrates the process of the anomaly detection task. First, the log message must be transformed into a format that is suitable for machine learning algorithms. Since anomaly detection models are often designed as neural networks [35, 7, 8] which we also utilize in this work, we consider the case of employing a neural network architecture. The output of a neural network is commonly transformed such that it can be interpreted as a probability or a probability distribution over one or multiple output neurons, from which further decisions can be inferred. Therefore, an anomaly detection method requires a final

decision on interpreting the neural network's output. From this, we conclude that an unsupervised anomaly detection method, which includes a neural network, consists of two parts: The *Anomaly Scoring* and the *Anomaly Decision*. The neural network learns the normal behavior in the first part and calculates an anomaly score for each log message. Therefore, the *Anomaly Scoring* must transform each log message into a real-valued output. The second part is the final anomaly detection decision, based on the anomaly scoring, which transforms the scores into a binary classification. Formally, the anomaly detection model  $\Phi$  for log data  $L$  is described by two functions,  $s$  and  $a$ , where  $s : L \rightarrow \mathbb{R}$  is the *Anomaly Scoring* function, represented by a neural network and  $a : \mathbb{R} \rightarrow \{0, 1\}$  is the *Anomaly Decision*. Therefore,  $\Phi = a(s(x))$ , where  $x$  is the preprocessed log data input.

#### 3.2. Problem Description

Two main challenges naturally arise when designing an unsupervised anomaly detection method, in our case on log data. On the one hand, the respective model is purely trained on normal log messages, which makes it difficult to identify abnormal log messages as such, as the computed anomaly scores cannot be interpreted, and thus scores of real abnormal log messages can have an arbitrary form. On the other hand, the decision boundary for separation of both classes is also configured based on normal log messages only. A solution to this problem needs to take these challenges into consideration.

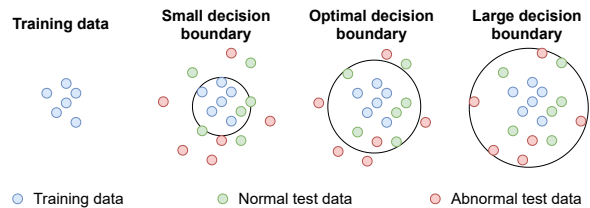


Figure 2. Illustration of the various possibilities for manually setting a decision boundary.

Figure 2 demonstrates the difficulty of setting a precise decision boundary by only utilizing normal data points from the training. The blue points represent normal data points, which are part of the training data. The green and red points represent normal and abnormal data points during the test phase. It can be observed that the green points are on average closer to the normal trained data points, yet considering them for configuration of the decision boundary would lead to a suboptimal separation. If the decision boundary is too small, too many log messages will be classified as false positive by the method. If the decision boundary is too

large, too many log messages will be classified as false negatives. Both scenarios result in weaker precision, recall, and F1 scores. This example demonstrates that both components, the *Anomaly Scoring* and the *Anomaly Decision*, have to be precise.

Due to the various factors influencing the training of a neural network, each trained anomaly scoring model will most likely produce different anomaly scores for the same data, and therefore the final decision boundary must be set individually for each model. Moreover, as the nature and relation of future normal and abnormal log messages cannot be known in advance, it is required that the decision boundary can cope with these uncertainties and still provide good decisions solely based on seen normal log messages.

## 4. Method

This section describes in detail our approach A2Log for unsupervised anomaly detection on log data, including our proposed solutions for anomaly scoring and anomaly decision. Prior to that, we briefly introduce the necessary preliminaries.

### 4.1. Preliminaries

Logging is commonly employed in order to investigate faulty behavior of systems and services and to increase dependability, which results in information being written to a log file. The log file documents the executions of the software and is created by log instructions (e.g. `printf()` or `log.info()`). Each log instruction results in a single log message, such that the complete log is a sequence of messages  $\mathcal{L} = (l_i : i = 1, 2, 3, \dots)$ . There is a commonly used separation in *meta-information* and *content*. The meta-information can contain various information, for example, timestamps or severity levels. The content is free text and consists of a static and a variable part. The static part is called log template. To access the content of a log message  $l_i$ , we write  $c_i$ .

In order to transform the content into a representation that an algorithm can process, methods from the research field of Natural Language Processing (NLP) are applied. Two main concepts are tokenization and embeddings.

**Tokenization.** This process splits written text into segments (e.g., words, word stems, or characters). The smallest indecomposable unit within a log content  $c_i$  is a token. Consequently, each log content  $c_i$  can be interpreted as a sequence of tokens  $t_i$

$$t_i = (w_j : w_j \in V, j = 1, 2, 3, \dots), \quad (1)$$

where  $w$  is a token,  $V$  is a set of all known tokens commonly referred to as the *vocabulary*,  $j$  is the positional index of a token within the token sequence. To access a specific token at position  $j$ , we write  $t_i^j$ . Thus, tokenization translates a text into a sequence of tokens using a vocabulary. The amount of tokens  $n_i$  for each log message and the structure of each token can vary, depending on the concrete tokenization method.

**Embeddings.** As the tokens themselves are words of the vocabulary  $V$ , they cannot be passed into a neural network directly. Furthermore, tokens do not provide any information about their similarity or difference to each other, hence, so called *embeddings* are used to compute a representation of the tokens such that a machine learning model can process it. Embeddings are real-valued vector representations  $\vec{v} \in \mathbb{R}^d$  of either token sequences or a single token; a transformation function  $g : V \rightarrow \mathbb{R}^d$  transforms a token  $w_j$  into an embedding  $\vec{v}_j$ . Thereby the same tokens receive the same embeddings. The sequence of embeddings  $\vec{e}_i$ , that is describing the corresponding sequence of tokens  $t_i$ , is defined as follows:

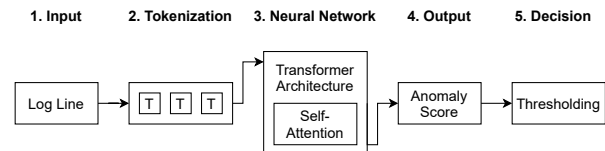
$$\vec{e}_i = (g(w_j) : w_j \in t_i, j = 1, 2, \dots, |t_i|), \quad (2)$$

To access the  $j$ -th embedding in a sequence of embeddings, we write  $\vec{e}_i(j)$ . Embeddings are trainable units adapted during the training process to represent the meaning of the underlying token or sequence of tokens.

### 4.2. A2Log

As a solution for an unsupervised anomaly detection method, we propose *A2Log*. *A2Log* consists of two parts: the *Anomaly Scoring* and the *Anomaly Decision*. To provide a general anomaly detection method, we only utilize the content of log messages, which is a commonality of different log types.

The schematic flow in Figure 3 of *A2Log* is as follows.



**Figure 3. Five steps to classify a log message utilizing a transformer architecture.**

During the first step, the content  $c_i$  of every log message  $l_i$  is extracted and in the process tokenized to a sequence of tokens  $t_i$  using the symbols `.`, `:` and whitespaces as separators. Subsequently, we further clean the resulting sequence of tokens

by replacing certain tokens with placeholders that adequately represent the original token without losing relevant information. We introduce placeholder tokens for hexadecimal values [HEX] or any number greater or equal 10 [NUM]. Finally, we prefix the sequence of transformed tokens with a special placeholder token [CLS] which will be beneficial later on. An exemplary log message

```
time.c: Detected 3591.142 MHz.
```

is thus transformed into a sequence of tokens  $t_i$ :

```
[[CLS], time, c, Detected, [NUM],  
[NUM], MHz].
```

The token sequence  $t_i$  serves as the input for our *Anomaly Scoring* model. Therefore, we utilize the encoder of the transformer architecture [32] with self-attention [27]. We have chosen this model architecture since it has already performed well in the domain of natural language processing. The encoder of the transformer architecture is applied to map sequences of tokens onto one d-dimensional vector (embedding), which is represented through the [CLS] token.

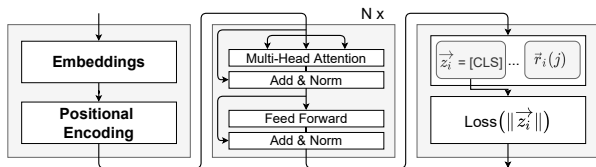


Figure 4. Transformer encoder architecture.

Figure 4 depicts the aforementioned network architecture. During the embedding step, for each token  $w_j$  in the token sequence  $t_i$ , an embedding  $\vec{e}_i(j)$  is obtained using the transformation function  $g$ . These embeddings are equal for the same tokens and modified during the training process. Since these sequences can vary in length, we truncate them to a fixed size  $u$  and fill up smaller sequences with padding tokens [PAD]. Hence, the employed architecture does not consider the order of the tokens, the input sequence is enriched with positional encoding [27, 36]. Based on this, the attention mechanism can take the order of the tokens into account.

The model then computes an output embedding for each truncated input embedding sequence, which summarizes the log message by utilizing the embeddings of all tokens. This output embedding is encoded in the embedding of the [CLS] token and also modified during training, by minimizing the loss. During the training process, the model is supposed to learn the meanings of the log messages, thereby getting an intuition of normal log messages. We denote the output of the model as  $\vec{z}_i = s(t_i)$  and use it throughout the remaining steps. Thereby the anomaly score is

calculated by the length of the output vector  $\|\vec{z}_i\|$ . These values must reflect the anomaly probability for each sequence. We set the scoring target for normal data to 0 as the absolute normal state, so that the likeliness of an anomaly in the log message is proportional to increasing positive values. The following objective function is utilized to ensure compact anomaly scores near to  $\|\vec{z}_i\| \geq 0$  for normal training data [37]:

$$\frac{1}{n} \sum_{i=1}^n (1 - y_i) \|\vec{z}_i\|^2 - y_i \cdot \log(1 - \exp(-\|\vec{z}_i\|^2)) \quad (3)$$

Thereby,  $y_i$  is the label for each input  $t_i$ , which implies that there must be two classes to train the neural network to avoid the implosion of the model. If there is only one class, the objective function will force the model to produce the same result independent of the input. Therefore, we define a stabilization class  $S$  as the second class which encompasses normal log messages but with different origins. This second class is constructed by utilizing normal log messages from other services, where we randomly sample an equal amount of log messages from each service. This gives the model an intuition of variety in log messages and, therefore, the model should be able to assign different anomaly scores to real anomalies based on log message characteristics. Eventually, the *Anomaly Decision* function assigns each log message an explicit label  $\in \{0, 1\}$  with

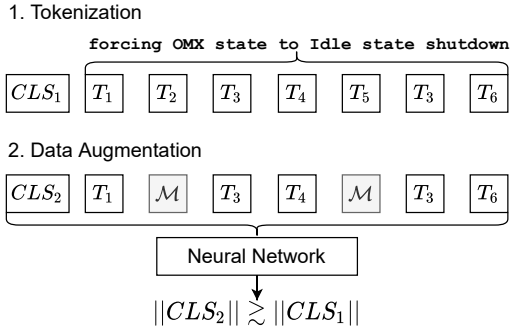
$$a(\|\vec{z}_i\|) = \begin{cases} 1, & \|\vec{z}_i\| > \epsilon, \\ 0, & \|\vec{z}_i\| \leq \epsilon \end{cases} \quad (4)$$

where  $\epsilon$  is a decision boundary that is used to obtain binary labels from anomaly scores.

As we do not know how the *Anomaly Score* will turn out for new normal log messages appearing in the future, we require a method that simulates deviations to the already trained log messages to understand how the model will react, and determines the decision boundary  $\epsilon$ . We simulate deviations to the normal training data by applying data augmentation.

Figure 5 depicts our method for data augmentation on log data. The data augmentation is applied to the log messages of the original training dataset  $l_{train}$  only, without the stabilization class  $S$ . First, the log messages are tokenized, followed by replacing  $\alpha \in \mathbb{N}$  tokens at random positions in each token sequence  $t_i$  with a masking token [M], described in Equation 5.

$$t'_i = (w_j : w_j \in V \vee w_j = [\mathcal{M}], j = 1, 2, \dots, |t_i|) \quad (5)$$



**Figure 5. Data augmentation with  $\alpha = 2$  replaced tokens.**

Then the trained model for the *Anomaly Scoring*  $s$  calculates a scalar value for each augmented token sequence  $t'_i$  which should be greater or equal than the scalar value for the original token sequence  $s(t'_i) \gtrsim s(t_i)$  because it contains  $\alpha$  unknown tokens  $[\mathcal{M}]$  and, therefore, does not comply with established knowledge. We gather all values - calculated by the model - for the augmented token sequences to receive a distribution  $D$ .

$$D = (s(t'_i) : t'_i \in l_{train}, 0 \leq i \leq |l_{train}|). \quad (6)$$

This distribution is the basis for the decision boundary  $\epsilon$ . We calculate the decision boundary by choosing the  $i$ -th percentile  $p$  of the distribution  $D$  and multiplying this value with a regulator variable  $\beta$ .

$$\epsilon(D, p, \beta) = D_p \cdot \beta \quad (7)$$

This formula calculates the final decision boundary from the given distribution for the deviated token sequences. The parameter  $p$  is a bias regulator, because there is a possibility, that the model calculates outliers for few normal augmented log messages. These outliers can occur when rare normal log messages resemble the stabilization class data and have high anomaly scores. With  $\beta$ , we can control whether we want to allow more deviations for normal data during the prediction phase.

## 5. Evaluation

To quantify the performance of A2Log, we test it on three real-world HPC log datasets and one industry dataset. Thereby, we compare our anomaly detection method with three other approaches and evaluate the performance of our anomaly decision boundary.

### 5.1. Experimental Setup

For the evaluation, we investigate four different training splits where we use the first 10%, 20%, 40%,

60% of normal data to train the anomaly scoring model and calculate the anomaly decision boundary. The remaining log messages are used for the testing. For each split, we examine the F1 score, precision, and recall. This setup is applied to the three publicly available datasets as well as to the industry dataset.

**Publicly Available Datasets.** We select three real-world datasets from HPC systems for evaluation as target systems, namely Blue Gene/L (BGL), Spirit, and Thunderbird (Tbird) [38]. From every dataset<sup>2</sup>, we utilize the first 5 million log messages.

To reveal the characteristics of the datasets, we have calculated the log templates for each dataset using Drain [39]. Table 1 depicts the count of normal and abnormal data, the count of log templates for the whole dataset, and the number of distinctive log templates that are present in the test dataset only, without being in the training dataset. It can be seen that especially for small amounts of training data, there are a lot of log messages in the test data, that are not seen in the training data, especially on the BGL dataset. That makes anomaly prediction difficult when only the normal data is trained, and the normal data in the test set deviates from it. Overall, 7 % to 15 % of the log lines are abnormal.

**Industry Dataset.** Furthermore, we investigate A2Log for an industry dataset that comes from the production environment of an IT service and cloud provider. Thereby the focus is to ensure a dependable storage service, by detecting anomalies in the log data from the underlying hardware. For this, we used log data from different disk controllers that are managing a variety of hard disks. The task is to identify anomalies by only training normal log data from some hard disks with no recorded anomalies and apply the model to new disks without manual optimization.

Table 2 shows the total number of training samples and the number of unique training samples, depending on the split. It also reveals the numbers of total and unique normal samples in the test data. Furthermore, it reveals the count of abnormal samples and the number of unique anomalies in the test dataset. It can be concluded that there are only very few different anomalies in the test dataset but they occur frequently. Furthermore, the rate of anomalies for the dataset is between 3 to 3.5 %, regarding the respective split. The dataset has 7,120,988 log lines in total.

**Benchmark Approaches.** In total, we evaluate four different methods, including A2Log. The first baseline is DeepLog, where we set the final decision boundary as best as possible. For the other three

<sup>2</sup><https://www.usenix.org/cfdr-data>

System	#normal	#anomalies	#templates	#unique log templates in test that do not appear in train for splits			
				10 %	20 %	40 %	60 %
BGL	4,399,503	348,460	1,571	1,318	1,232	1,158	1,077
Thunderbird	4,773,713	226,287	1,302	232	205	128	61
Spirit	4,235,110	764,890	1,457	1,091	1,028	297	129

Table 1. Dataset description of BGL, Thunderbird and Spirit.

Train split	#train	#unique train	#test	#unique test	#unique test abnormal	#test abnormal
0.1	711,841	226,375	6,409,147	1,270,089	28	196,534
0.2	1,424,198	465,666	5,696,790	106,0947	27	196,519
0.4	2,848,395	536,947	4,272,593	987,417	23	146,099
0.6	4,272,593	895,520	2,848,395	614,089	18	100,827

Table 2. Dataset description of the industry dataset.

methods, we train our described transformer model for the anomaly scoring and then calculate the decision boundary differently. First, we calculate the decision boundary based on the 3-sigma method. The 3-sigma decision boundary is calculated on the anomaly scores of the trained model for all training data, excluding the stabilization class. The second benchmark *Best* is the best possible result for the transformer model, which is achieved through utilizing the test data, including the abnormal samples, to calculate the optimal decision boundary. Therefore, this benchmark cannot be exceeded and is considered as an absolute upper limit. Therefore, the goal is to get as close as possible to this result. We then evaluate *A2Log* against these three benchmarks.

## 5.2. Implementation Details

First, we tokenize the content of each log line as described in section 4 and then truncate all sequences of tokens to the same length  $|t_i| = 20$ . We set the dimensionality  $d$  of the embeddings to 128. The weights of the embeddings are initialized with *Xavier* [40]. For the anomaly scoring model training, we use a hidden size of 256, a batch size of 1024, a dropout rate of  $5 \cdot 10^{-2}$ . For the optimization task, we use the Adam optimizer with a learning rate of  $10^{-4}$  and a weight decay of  $5 \cdot 10^{-5}$  in every experiment.

Since we need a stabilization class  $S$  to train our model, we use 60k random messages from other datasets. The corresponding mappings are displayed in Table 3.

For the three datasets (BGL, Thunderbird, and Spirit), we utilize 120,000 log lines each as a

Dataset for evaluation	#logs for stabilization from			#total
	BGL	Tbird	Spirit	
BGL	/	60000	60000	120000
Tbird	60000	/	60000	120000
Spirit	60000	60000	/	120000
Industry dataset	60000	60000	60000	180000

Table 3. Log messages utilized for the stabilization class.

stabilization class. For the industry dataset, we utilize a total of 180,000 log lines as the stabilization class. Due to the disparity in the number of training samples of the normal log messages and the stabilization class, we use a weighted sampler for the training to balance the training. All models are trained until an average loss of 0.01 per sample, or a maximum of 50 epochs, is reached.

The anomaly decision function  $a$  is parameterized as shown in Table 4.

Parameter	BGL	Thunderbird	Spirit	industry dataset
$\alpha$	1	1	1	1
$p$	0.95	0.95	0.95	0.95
$\beta$	2.5	2.5	5.0	2.0

Table 4. Decision boundary parameters to calculate  $\epsilon$ .

We set  $\alpha$  to 1 in each experiment for each dataset to replace only one token of the log message with the unknown token for the data augmentation. Likewise, in each experiment, we set  $p$  to 0.95 to filter out 5 % of the augmented data with the highest anomaly scores.

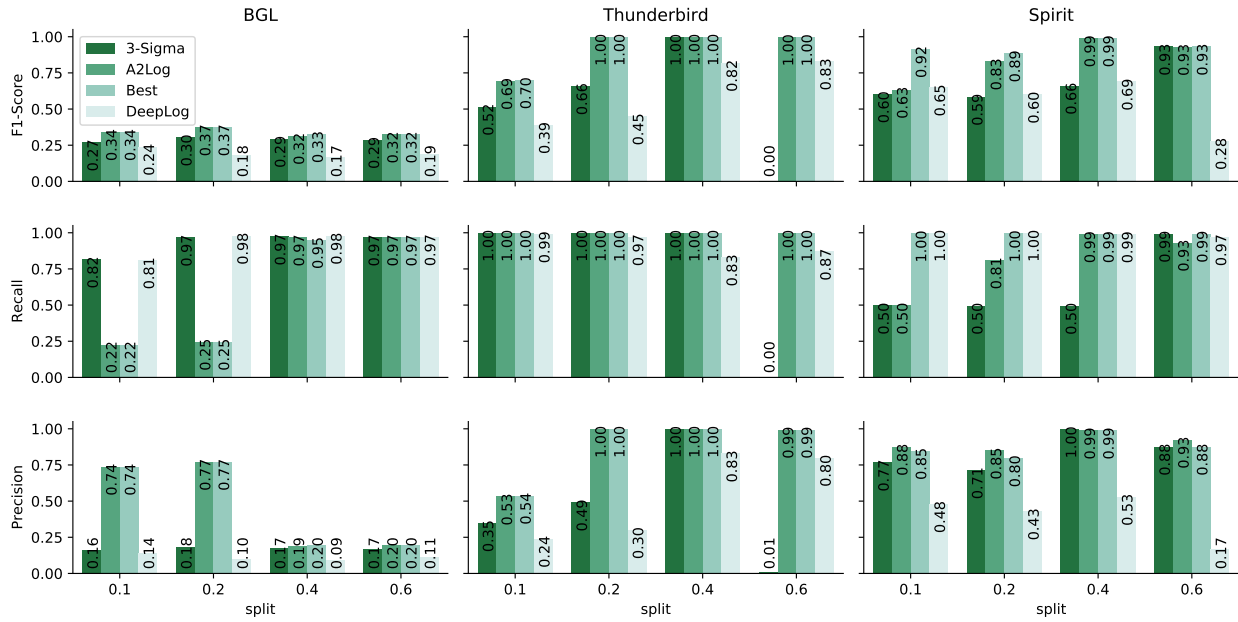


Figure 6. Evaluation on BGL, thunderbird and spirit.

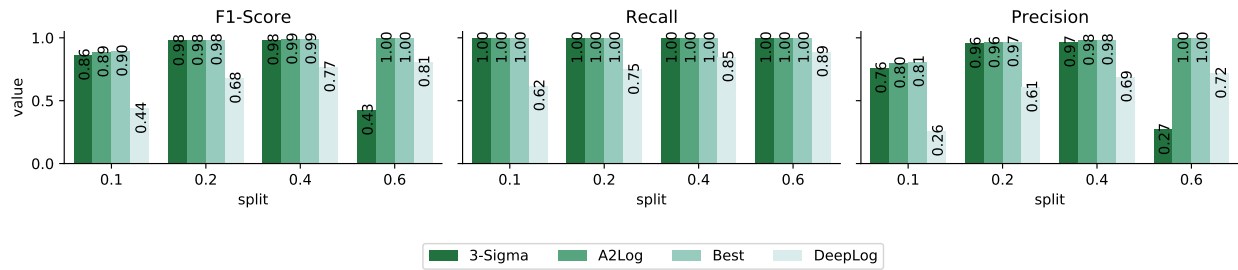


Figure 7. Evaluation on the industry dataset.

Furthermore, we adjust  $\beta$  for each dataset but keep it the same for each experiment on the respective dataset.

### 5.3. Results

For all evaluations, we do every experiment three times and depict the best results in terms of the F1-score in Figure 6 for the three publicly available datasets and in Figure 7 for the industry dataset.

Figure 6 depicts the results of *Best*, 3-Sigma and Deeplog, compared with *A2Log*. It shows the respective approaches for different amounts of training data of 10%, 20%, 40%, 60% of the respective dataset. It can be seen that *A2Log* is superior or equal to DeepLog and the baseline 3-Sigma in all experiments. In addition, *A2Log* achieves almost the same F1 scores as *Best* in all experiments. A noticeable aspect of the Spirit dataset is that as soon as only a few learning samples are available, the decision boundary function of *A2Log*

no longer turns out to be that precise. Nevertheless, the other training splits for the Spirit dataset show that our unsupervised *A2Log* approach can perform equally to an optimal decision boundary. The results on the BGL dataset are not particularly high, as there is a fundamental concept drift in the data, which can be seen by the fact that even for high training splits of 60 % there are still many log templates present in the test dataset that are not present in the training dataset as shown in Table 1.

In addition, the two parameters  $\alpha$  and  $p$  can be set to 1 and 0.95 independently of the dataset. This shows that the method is robust in parameter choice and independent of the dataset.

Figure 7 depicts the performance of *A2Log* for different training splits. Thereby it compares *A2Log* with the best possible decision boundary and the 3-sigma boundary. It can be observed that *A2Log* outperforms 3-Sigma and is as good as the best possible



decision boundary. When *A2Log* can train enough data, it is able to identify anomalies on new hard disks with perfect scores.

## 6. Conclusion

Anomaly detection methods have become increasingly important to ensure the dependable and stable operation of IT services, including their serviceability. However, existing unsupervised anomaly detection methods are applied under constrained assumptions for the final anomaly decision. Therefore we propose *A2Log*, to address the current limitations of unsupervised anomaly detection methods. Unlike other unsupervised approaches, it calculates its decision boundary for the final decision by exploring the model behavior based on augmented training data. With data augmentation, we simulate deviations in log data that occur from service updates over time. We evaluate our approach on three publicly available datasets and one industry dataset. Thereby, we show the effectiveness of our approach when it is applied in an unsupervised setting.

Even if we can simulate deviations of the normal log lines with the help of our data augmentation and thus classify new normal log lines as such, we see as a limitation that a fundamental concept drift in the data pushes the method to its limits and thus leads to misclassifications. Nevertheless, *A2Log* can keep up with the optimal decision boundary, which can only be calculated by utilizing available anomaly examples. Furthermore, we are able to show that the entire method outperforms *DeepLog* and *3-Sigma* as a baseline decision boundary.

As further work, we consider additional adjustments to the data augmentation process, to integrate it into other anomaly detection methods. Thereby, we want to apply data augmentation techniques that are well known in the research area of natural language processing. Furthermore, we want to extend *A2Log* to other anomaly detection settings, e.g. supervised or weakly supervised settings.

## References

- [1] D. Rosendo, G. Leoni, D. Gomes, A. Moreira, G. Gonçalves, P. Endo, J. Kelner, D. Sadok, and M. Mahloo, "How to improve cloud services availability? investigating the impact of power and it subsystems failures," in *Proceedings of the 51st Hawaii international conference on system sciences*, 2018.
- [2] G. L. Santos, P. T. Endo, G. Gonçalves, D. Rosendo, D. Gomes, J. Kelner, D. Sadok, and M. Mahloo, "Analyzing the it subsystem failure impact on availability of cloud services," in *2017 IEEE symposium on computers and communications (ISCC)*, pp. 717–723, IEEE, 2017.
- [3] L. Baier, N. Kühn, and G. Satzger, "How to cope with change?-preserving validity of predictive services over time," in *Proceedings of the 52nd Hawaii International Conference on System Sciences*, 2019.
- [4] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM computing surveys (CSUR)*, vol. 41, no. 3, pp. 1–58, 2009.
- [5] T. Wittkopp and A. Acker, "Decentralized federated learning preserves model and data privacy," in *International Conference on Service-Oriented Computing*, pp. 176–187, Springer, 2020.
- [6] H. Guo, S. Yuan, and X. Wu, "Logbert: Log anomaly detection via bert," *arXiv preprint arXiv:2103.04475*, 2021.
- [7] M. Du, F. Li, G. Zheng, and V. Srikumar, "Deeplog: Anomaly detection and diagnosis from system logs through deep learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1285–1298, 2017.
- [8] S. Nedelkoski, J. Bogatinovski, A. Acker, J. Cardoso, and O. Kao, "Self-supervised log parsing," *arXiv preprint arXiv:2003.07905*, 2020.
- [9] P. Liu, X. Wang, C. Xiang, and W. Meng, "A survey of text data augmentation," in *2020 International Conference on Computer Communication and Network Security (CCNS)*, pp. 191–195, IEEE, 2020.
- [10] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 1, pp. 1–48, 2019.
- [11] G.-J. Qi and J. Luo, "Small data challenges in big data era: A survey of recent progress on unsupervised and semi-supervised methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [12] S. He, J. Zhu, P. He, and M. R. Lyu, "Experience report: System log analysis for anomaly detection," in *ISSRE*, IEEE, 2016.
- [13] Y. Liang, Y. Zhang, H. Xiong, and R. Sahoo, "Failure prediction in ibm bluegene/l event logs," in *ICDM*, IEEE, 2007.
- [14] I. Jolliffe, "Principal component analysis," *Encyclopedia of statistics in behavioral science*, 2005.
- [15] J. Breier and J. Branišová, "Anomaly detection from log files using data mining techniques," in *Information Science and Applications*, pp. 449–457, Springer, 2015.
- [16] M. Chen, A. X. Zheng, J. Lloyd, M. I. Jordan, and E. Brewer, "Failure diagnosis using decision trees," in *ICAC*, IEEE, 2004.
- [17] W. Xu, L. Huang, A. Fox, D. Patterson, and M. I. Jordan, "Detecting large-scale system problems by mining console logs," in *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*, pp. 117–132, 2009.
- [18] J.-G. Lou, Q. Fu, S. Yang, Y. Xu, and J. Li, "Mining invariants from console logs for system problem detection.," in *USENIX Annual Technical Conference*, pp. 1–14, 2010.
- [19] M. Cinque, D. Cotroneo, and A. Pecchia, "Event logs for the analysis of software failures: A rule-based approach," *IEEE Transactions on Software Engineering*, vol. 39, no. 6, pp. 806–821, 2012.

- [20] E. Baseman, S. Blanchard, Z. Li, and S. Fu, "Relational synthesis of text and numeric data for anomaly detection on computing system logs," in *2016 15th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pp. 882–885, IEEE, 2016.
- [21] M. Landauer, M. Wurzenberger, F. Skopik, G. Settanni, and P. Filzmoser, "Dynamic log file analysis: An unsupervised cluster evolution approach for anomaly detection," *Computers & Security*, vol. 79, pp. 94–116, 2018.
- [22] X. Zhang, Y. Xu, Q. Lin, B. Qiao, H. Zhang, Y. Dang, C. Xie, X. Yang, Q. Cheng, Z. Li, *et al.*, "Robust log-based anomaly detection on unstable log data," in *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pp. 807–817, 2019.
- [23] A. Brown, A. Tuor, B. Hutchinson, and N. Nichols, "Recurrent neural network attention mechanisms for interpretable system log anomaly detection," in *Proceedings of the First Workshop on Machine Learning for Computing Systems*, pp. 1–8, 2018.
- [24] K. Yin, M. Yan, L. Xu, Z. Xu, Z. Li, D. Yang, and X. Zhang, "Improving log-based anomaly detection with component-aware analysis," in *2020 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pp. 667–671, IEEE, 2020.
- [25] M. Nicolau, J. McDermott, *et al.*, "A hybrid autoencoder and density estimation model for anomaly detection," in *International Conference on Parallel Problem Solving from Nature*, pp. 717–726, Springer, 2016.
- [26] M. Sakurada and T. Yairi, "Anomaly detection using autoencoders with nonlinear dimensionality reduction," in *Proceedings of the MLSDA 2014 2nd workshop on machine learning for sensory data analysis*, pp. 4–11, 2014.
- [27] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *arXiv preprint arXiv:1706.03762*, 2017.
- [28] S. He, J. Zhu, P. He, and M. R. Lyu, "Experience report: System log analysis for anomaly detection," in *2016 IEEE 27th International Symposium on Software Reliability Engineering (ISSRE)*, pp. 207–218, 2016.
- [29] R. Yang, D. Qu, Y. Gao, Y. Qian, and Y. Tang, "Nlsalog: An anomaly detection framework for log sequence in security management," *IEEE Access*, vol. 7, 2019.
- [30] W. Meng, Y. Liu, Y. Zhu, S. Zhang, D. Pei, Y. Liu, Y. Chen, R. Zhang, S. Tao, P. Sun, *et al.*, "Loganomaly: Unsupervised detection of sequential and quantitative anomalies in unstructured logs," in *IJCAI*, vol. 7, pp. 4739–4745, 2019.
- [31] L. Yang, J. Chen, Z. Wang, W. Wang, J. Jiang, X. Dong, and W. Zhang, "Semi-supervised log-based anomaly detection via probabilistic label estimation," in *ICSE*, IEEE, 2021.
- [32] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [33] I. Steinwart, D. Hush, and C. Scovel, "A classification framework for anomaly detection.," *Journal of Machine Learning Research*, vol. 6, no. 2, 2005.
- [34] A. Ramponi and B. Plank, "Neural unsupervised domain adaptation in nlp—a survey," *arXiv preprint arXiv:2006.00632*, 2020.
- [35] S. Naseer, Y. Saleem, S. Khalid, M. K. Bashir, J. Han, M. M. Iqbal, and K. Han, "Enhanced network anomaly detection based on deep neural networks," *IEEE access*, vol. 6, pp. 48231–48246, 2018.
- [36] J. Gehring, M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin, "Convolutional sequence to sequence learning," in *ICML* (D. Precup and Y. W. Teh, eds.), PMLR, 2017.
- [37] S. Nedelkoski, J. Bogatinovski, A. Acker, J. Cardoso, and O. Kao, "Self-attentive classification-based anomaly detection in unstructured logs," in *2020 IEEE International Conference on Data Mining (ICDM)*, pp. 1196–1201, 2020.
- [38] A. Oliner and J. Stearley, "What supercomputers say: A study of five system logs," in *37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07)*, pp. 575–584, IEEE, 2007.
- [39] P. He, J. Zhu, Z. Zheng, and M. R. Lyu, "Drain: An online log parsing approach with fixed depth tree," in *2017 IEEE international conference on web services (ICWS)*, pp. 33–40, IEEE, 2017.
- [40] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256, JMLR Workshop and Conference Proceedings, 2010.