

# A Comparison of LSTM and XGBoost for Predicting Firemen Interventions <sup>★</sup>

Selene Cerna<sup>1</sup>[0000-0003-1690-1279], Christophe Guyeux<sup>2</sup>[0000-0003-0195-4378], Héber H. Arcolezi<sup>2</sup>[0000-0001-8059-7094], Raphaël Couturier<sup>2</sup>[0000-0003-1490-9592], and Guillaume Royer<sup>3</sup>[0000-0000-0000-0000]

<sup>1</sup> São Paulo State University (UNESP), Ilha Solteira-SP, Brazil  
`selene.cerna@unesp.br`

<sup>2</sup> Femto-ST Institute, UMR 6174 CNRS, Univ. Bourgogne Franche-Comté, France  
`{christophe.guyeux,heber.hwang-arcolezi,raphael.couturier}@univ-fcomte.fr`

<sup>3</sup> SDIS 25, Besançon, France  
`guillaume.royer@sdis25.fr`

**Abstract.** In several areas of the world such as France, fire brigades are facing a constant increase in the number of their commitments, some of the main reasons are related to the growth and aging of the population and others to global warming. This increase occurs principally in constant human and material resources, due to the financial crisis and the disengagement of the states. Therefore, forecasting the number of future interventions will have a great impact on optimizing the number and the type of on-call firefighters, making it possible to avoid too few firefighters available during peak load or an oversized guard during off-peak periods. These predictions are viable, given firefighters' labor is conditioned by human activity in general, itself correlated to meteorological data, calendars, etc. This article aims to show that machine learning tools are mature enough at present to allow useful predictions considering rare events such as natural disasters. The tools chosen are XGBoost and LSTM, two of the best currently available approaches, in which the basic experts are decision trees and neurons. Thereby, it seemed appropriate to compare them to determine if they can forecast the firefighters' response load and if so, if the results obtained are comparable. The entire process is detailed, from data collection to the predictions. The results obtained prove that such a quality prediction is entirely feasible and could still be improved by other techniques such as hyperparameter optimization.

**Keywords:** Long short-term memory · Extreme gradient boosting · Firemen interventions · Machine learning · Forecasting.

---

<sup>★</sup> This work was supported by the EIPHI Graduate School (contract "ANR-17-EURE-0002"), by the Region of Bourgogne Franche-Comté CADRAN Project, by the Interreg RESponse project, and by the SDIS25 firemen brigade. We also thank the supercomputer facilities of the Mésocentre de calcul de Franche-Comté

## 1 Introduction

Most of the problems faced on a day-to-day basis by fire brigades are related to the increase in the number of interventions over time and with the management of an insufficient budget. This results in personnel and equipment shortages and affecting the response time to the incidents. Therefore, taking advantage of their data gathered through the years to build models that can predict the occurrence of an intervention in the future would help in establishing better strategies to nurse the community and reduce the response time. Consequently, more lives would be saved with fewer efforts.

Making a review of the literature, researches on the specific field of forecasting the number, type or location of interventions for fire departments is still scarce in the literature [5]. For this reason, the present work compares the use of two machine learning (ML) methods: the Extreme Gradient Boosting (XGBoost), which is based on decision trees that highly optimize the processing time and model complexity. And, the Long Short-Term Memory (LSTM), a highlighting variation of the Recurrent Neural Network (RNN) and introduced by [12], which has shown a remarkable performance in sequential data applications along with overcoming the vanishing gradient problem presented in the RNN [10, 11]. With the primary objective of providing a data-driven decision-making approach for fire departments to forecast the number of interventions in the next hour. As references, it was taken some researches with LSTM: short-term traffic speed and flow predictions [13, 8] and a survey on the analysis of eight LSTM variants on three tasks: speech recognition, handwriting recognition, and polyphonic music modeling [11]. For XGBoost one can find researches predicting traffic flow prediction using ensemble decision trees for regression [4] and with a hybrid deep learning framework [15].

The following sections of this paper are structured as: in Section 2.1 the way the data were acquired and encoded is presented; in Section 2.2 a short description of LSTM and XGBoost methods is provided; in Section 3 prediction results are described and a discussion to highlight the results is made, and in Section 4 concluding thoughts and forthcoming works are given.

## 2 Materials and Methods

### 2.1 Data Acquisition and Encoding

It was considered two sources. The main contains information about all the interventions recorded from 2006 to 2017 by the fire and rescue department SDIS25, in the region of Doubs-France. And the second contains external variables such as weather, traffic, holidays, etc.

First, the date and the time of each intervention were extracted, in order to recognize time patterns on the occurrence of incidents. For example, it was noticed that there were more interventions occurring during the day. Besides, meteorological data were considered, which contributes significantly to the forecast of the number of incidents (e.g., road accidents are related to road surface

condition). Holidays and academic vacations were also taken into account in view of young people tend to go out during these periods. Thus, a dictionary was organized as it is described in [5], section III, subsection A, with the following differences:

- The height of the six most important rivers of the Doubs department was considered. The average, the standard deviation and the number of readings belonging to the block of 1h were used [2].
- From the Skyfield library [3], it was taken the distance between the Earth and the Moon to examine its influence in natural disasters.
- Festivities such as Ramadan, Eurockéennes, Percée du Vin Jaune and the FIMU were included as indicators with values 1 for the eve, duration and a day after, and 0 for normal days.
- After analyzing the data, it was discovered that leap years have an impact on the variable of the day in the year. For instance, July 14th (the National Day of France) is not the same day when the month of February has 29 days. For this reason, February 29th of 2008, 2012 and 2016 were removed.

The data were transformed into our learning format employing two methods from Scikit-learn library [14]. The “StandardScaler” method was applied to numerical variables such as year, hour, wind speed and direction, humidity, nebulosity, dew point, precipitations, bursts, temperature, visibility, chickenpox, influenza, and acute diarrhea statistics, rivers height and moon distance; which re-scales the distribution of values to zero mean and unit variance. The “OneHotEncoder” method was employed to convert into indicators categorical variables such as Bison Futé’s values, day, day of the week, day of the year and month, holidays, barometric trend and festivities. The original target values were kept (the number of interventions) because the distribution of the interventions count is better represented by discrete values.

The organization of each sample consisted of joining the extracted features with the number of interventions of the previous 169 hours (1 week plus 1h). Eventually, the data set is considered as sequential data and converted to supervised learning, i.e., the target is the number of interventions in the next hour ( $t + 1$ ) of a present sample ( $t$ ).

## 2.2 Machine Learning Techniques for Predicting Firefighters Interventions

**Long Short-Term Memory** Its memory cell consists of one principal layer and three gate controllers: input, forget and output. The principal layer analyzes the present entry  $x_t$  and the preceding short-term state  $h_{t-1}$ . The input gate regulates the flow of new memories. The forget gate controls which memories will be eliminated from the previous long-term state  $c_{t-1}$ , and with the new memories it is obtained the new long-term state  $c_t$ . The output gate establishes which memories will be considered as the new output of the LSTM cell for a specific time step, i.e., the  $y(t)$ , which at some point during the operation is equal

to the new short-term state  $h(t)$  [9]. The process is mathematically expressed as:

$$i_t = \sigma(W_{xi}^T \cdot x_t + W_{hi}^T \cdot h_{t-1} + b_i) \quad (1)$$

$$f_t = \sigma(W_{xf}^T \cdot x_t + W_{hf}^T \cdot h_{t-1} + b_f) \quad (2)$$

$$o_t = \sigma(W_{xo}^T \cdot x_t + W_{ho}^T \cdot h_{t-1} + b_o) \quad (3)$$

$$g_t = \tanh(W_{xg}^T \cdot x_t + W_{hg}^T \cdot h_{t-1} + b_g) \quad (4)$$

$$c_t = f_t \otimes c_{t-1} + i_t \otimes g_t \quad (5)$$

$$y_t = h_t = o_t \otimes \tanh(c_t) \quad (6)$$

where  $W_{xi}$ ,  $W_{xf}$ ,  $W_{xo}$  and  $W_{xg}$  are the weight matrices for their connection to the input vector  $x_t$ ;  $W_{hi}$ ,  $W_{hf}$ ,  $W_{ho}$  and  $W_{hg}$  are the weight matrices for their connection to the previous short-term state  $h_{t-1}$ ; and  $b_f$ ,  $b_g$ ,  $b_i$  and  $b_o$  are the bias terms of each layer. For more details about LSTM, see [12, 10, 9].

Our LSTM model was developed with Keras library [7]. It was built with one LSTM layer and 6000 neurons, one dense layer with one neuron as output and a last layer with the LeakyReLU activation function, considering 0.1 in the negative slope coefficient. The time step was one per input. For the training phase, the Stochastic Gradient Descent optimizer was used with a learning rate of 0.01, momentum and decay values of 0.0001, Poisson as loss function, a batch size of 64 and 200 epochs with an ‘‘EarlyStopping’’ of 10 epochs to monitor the loss function decrease of the validation set.

**Extreme Gradient Boosting** XGBoost uses a new regularization approach over the conventional Gradient Boosting Machines (GBMs) to significantly decrease the complexity. In order to measure the performance of a model given a certain data set, XGBoost defines an objective function considering the training loss  $L(\theta)$  and regularization  $\Omega(\theta)$  terms, where the latter penalizes the complexity of the model and prevents the overfitting, and  $\theta$  refers to the parameters that will be discovered during the training (Eq. 7). The result model  $\hat{y}_i^{(t)}$  at training the round  $t$  is the combination of  $k$  trees, i.e., an additive strategy is applied during the training, one new tree that optimizes the system  $f_t(x_i)$  is added at a time to the model  $\hat{y}_i^{(t-1)}$  generated in the previous round, where  $x_i$  is the input (Eq. 8). To determine the complexity of the tree  $\Omega(f)$ , [6] proposed an approach that defines it as Eq. 9, where the first term  $\gamma T$  evaluates the number of leaves  $T$ , taking  $\gamma$  as a constant, and the second term computes  $L2$  norm of leaves scores  $w_j$ . In Eq. 10 and Eq. 11,  $g_i$  and  $h_i$  respectively, are the first and second order partial derivatives after taking the Taylor expansion of the loss function

chosen,  $I_j = \{i|q(x_i) = j\}$  is the group of indices of data points attributed to the  $j$ -th leaf and  $q(x)$  is the structure of the tree. Finally, in the objective function, the argument of the minimum and the minimum of the quadratic function for the single variable  $w_j$  are taken, considering  $q(x)$  as fixed and  $\lambda$  as a very small constant value, the outcomes are Eq. 12 and Eq. 13, where the latter assesses the quality of a tree structure, i.e., an smaller score is better [6].

$$obj(\theta) = L(\theta) + \Omega(\theta) \quad (7)$$

$$\hat{y}_i^{(t)} = \sum_{k=1}^t f_k(x_i) = \hat{y}_i^{(t-1)} + f_t(x_i) \quad (8)$$

$$\Omega(f) = \gamma T + \frac{1}{2} \lambda \sum_{j=1}^T w_j^2 \quad (9)$$

$$G_j = \sum_{i \in I_j} g_i \quad (10)$$

$$H_j = \sum_{i \in I_j} h_i \quad (11)$$

$$w_j^* = -\frac{G_j}{H_j + \lambda} \quad (12)$$

$$obj^* = -\frac{1}{2} \sum_{j=1}^T \frac{G_j^2}{H_j + \lambda} + \gamma T \quad (13)$$

Our XGBoost model was improved using a GridSearchCV procedure from the Scikit-learn library [14]. The best model used in this research has a max depth of 3, a learning rate of 0.1, the learning task is Count and the learning objective is Poisson, which is for data counting problems; the remaining parameters were kept as default.

### 3 Prediction Results and Discussion

#### 3.1 Prediction Results

The metrics defined to evaluate the results are the Root Mean Square Error (RMSE) and the Mean Absolute Error (MAE). And because the events to be predicted are countable, the accuracy score was also considered with a margin of error zero (ACC0E) that represents the number of exact predictions reached, with a margin of error less or equal to one (ACC1E) and to two (ACC2E), which provides feasible results for real applications. In order to discover unusual years through the analysis of the prediction metrics, during each iteration, each year is predicted (it is considered as testing set), the remaining years are used as training and validation sets (e.g., to predict 2006, 2007-2017 were used as learning sets;

to predict 2007, 2006 and 2008-2017 were used as learning set, etc). Naturally, this is not a real case, but it provides information about how well can each year be predicted and why some years present atypical results. The data set was not cleaned from possible outliers, such as natural disasters (e.g., storms, fires, floods) and strikes that were found in our search analysis. Considering that in real-world applications, the system must perform well in such conditions it is worth maintaining these occurrences and evaluates the performances of the proposed methods.

Thus, Table 1 presents data analysis of the interventions per year, with the metrics: total number of interventions (Total Interv.), the average (Average), the standard deviation (Std. Dev.) and the maximum number of interventions (Max. Interv.).w Table 2 presents the results of the forecast to both LSTM and XGBoost models for all years (2006-2017). Fig. 1 represents the total number of interventions per year. Fig. 2, Fig. 3 and Fig. 4 illustrates responses of LSTM and XGBoost models on 100 samples trying to predict an unusual number of interventions occurred in 2010, 2011 and 2016 as a result of natural disasters in the Franche-Comté region. Moreover, forecasting results on 100 samples for 2017 are presented in Fig. 5, which is the year that only considers past years in the training process and presents an uncommon behavior due to ambulance strikes and climate conditions presented in the Doubs region during that year. Lastly, taking into account that the LSTM NN and the XGBoost models predict real values (e.g., 5.67 interventions), results were transformed to the closest integer (e.g., 6 interventions) for being coherent with real-world applications.

### 3.2 Discussion

The purpose of this research was to develop and evaluate two ML methods on forecasting the number of future firefighters interventions using data from 2006 to 2017, dividing them in training, validation and testing sets. As presented in Table 2, one can see that with reasonable efforts on features and relatively basic use of the XGBoost and LSTM techniques, quite good predictions results were obtained. Furthermore, it was noted that the results in both methods were very

**Table 1.** Data analysis of the interventions during 2006-2017

Year	Total Interv.	Average	Std. Dev.	Max. Interv.
2006	17,375	1.98	2.04	30
2007	19,368	2.21	2.06	13
2008	18,037	2.05	1.95	16
2009	28,719	3.27	3.34	84
2010	29,656	3.38	3.05	93
2011	33,715	3.84	3.66	48
2012	29,070	3.31	2.50	26
2013	29,830	3.40	2.48	30
2014	30,689	3.50	2.55	22
2015	33,586	3.83	2.68	21
2016	34,434	3.93	3.13	85
2017	37,674	4.30	2.94	22




Fig. 1. Total number of interventions per year.

Table 2. Prediction results on data 2006-2017

Year	LSTM					XGBoost				
	RMSE	MAE	ACC0E	ACC1E	ACC2E	RMSE	MAE	ACC0E	ACC1E	ACC2E
2006	1.60	1.13	28.28%	73.04%	90.99%	1.61	1.16	25.55%	73.27%	90.86%
2007	1.63	1.19	27.27%	70.91%	89.06%	1.66	1.20	26.19%	71.48%	88.83%
2008	1.59	1.16	26.83%	71.68%	90.28%	1.64	1.22	24.45%	69.94%	89.55%
2009	2.28	1.49	22.72%	62.28%	83.00%	2.39	1.58	21.59%	59.04%	80.36%
2010	2.32	1.49	23.17%	61.96%	81.92%	2.22	1.51	22.65%	60.82%	81.50%
2011	2.49	1.68	21.05%	57.54%	78.92%	2.55	1.69	21.07%	58.16%	78.93%
2012	2.06	1.53	21.30%	58.26%	81.11%	2.08	1.55	21.16%	58.03%	80.02%
2013	2.05	1.53	21.15%	58.81%	80.58%	2.06	1.54	20.91%	58.68%	80.22%
2014	2.04	1.52	21.26%	59.10%	81.17%	2.06	1.52	21.47%	59.37%	81.00%
2015	2.09	1.58	21.14%	56.41%	79.49%	2.09	1.56	21.51%	57.70%	79.48%
2016	2.64	1.71	18.94%	53.91%	77.51%	2.58	1.67	19.16%	55.49%	78.42%
2017	2.26	1.69	19.90%	54.63%	76.80%	2.27	1.68	20.38%	55.59%	76.94%

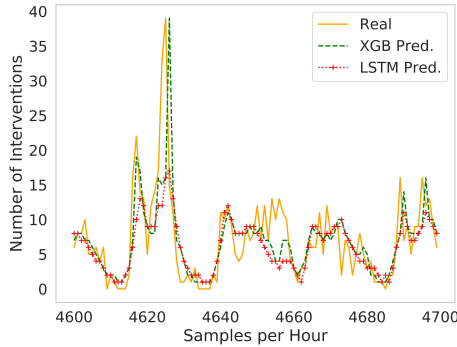


Fig. 2. Predictions for 2010.


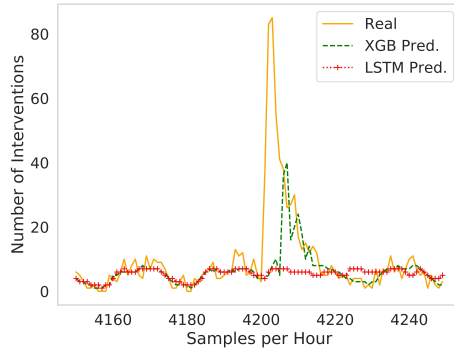



Fig. 3. Predictions for 2011.

similar, one for which the basic expert is a neuron (LSTM) and the other a decision tree (XGBoost).



**Fig. 4.** Predictions for 2016.



**Fig. 5.** Predictions for 2017.

Moreover, as one can see in Figs. 2-4, the XGBoost technique is a little more robust to outlier data than LSTM, where the first recognized the peak occurrences better during natural disasters. Considering that these occurrences are highly likely to happen in the future and fire brigades pursue to nurse its community better, real systems must be prepared to face the input data with uncommon values. Notwithstanding, the LSTM model presented better metrics values and accuracies for almost all of the years, which represents that in normal conditions and even with higher error values during peak occurrences, its metrics outstand those from XGBoost. Moreover, the use of deeper layers and more time steps could improve results by better generalizing the data.

Additionally, as presented in Table 1 and in Fig. 1, an increment in the number of interventions throughout the years is clearly highlighted, which could be probably due to population-aging and growth. However, one can notice an abnormal increment from 2008 to the years 2009 and 2011, in which natural disasters took place, i.e., in contrast to the aforementioned years, 2012-2015 follow a regular pattern of increment. This characteristic is also noted analyzing the metrics average, where from approximately 2 interventions per hour in 2008 increases to almost 4 in 2011; the standard deviation with higher values are for the years 2009-2011 and 2016, where the data were more sparse due to peak value occurrences during natural disasters, which is also well represented by the maximum number of interventions.

Also, in Table 2 one can observe a high increment in the RMSE and MAE metrics and a decrement in the ACC0E, ACC1E and ACC2E metrics starting from 2009. Initially, for the years 2009-2011 and 2016, poor metrics results are obtained probably occasioned by the outlier data. However, for 2012-2015 this increment is also following a normal pattern comparing to the increment of the total number of interventions. Finally, 2017, which is the most realistic prediction, presents lower ACC02E accuracy, which is probably because of the typical increment over the years and because of some factors that could not be detected by the models as outlier data, i.e., the increment was not just for a few hours like peaks (e.g., the Max. Interv. is just 22), but for many samples. In our search



analysis, we found that there was an ambulance strike that lasted 29 days between September and October, resulting in more incident attendance for the fire brigade. Also, it was found online sources related to an increment of 60% in the number of interventions for the Doubs department caused by a heatwave that occurred in June [1].

Therefore, for normal years, i.e., without outlier data the proposed models could achieve a good prediction, e.g., for 2006-2008 both models could predict with a high level of accuracy considering the ACC1E and ACC2E metrics approximately 73% and 90% were accomplished respectively. And, for years in which social or natural circumstances affected directly the prediction results, the scores are still acceptable for practical purposes, i.e., we do recognize that an intelligent system with accuracy between 50% and 70%, could not be used as first decision-making approach for fire brigades. However, we believe that results can even be improved by adding significant features and developing new models.

## 4 Conclusion

The development of intelligent systems to predict the number of interventions at a given time into the future could help fire brigades around the world to efficiently prepare themselves for future incidents. This paper presented two well-known machine learning methods, the LSTM and the XGBoost to predict the number of interventions for the next hour that firefighters would face in the region of Doubs-France. To validate the performance of both methods, a data set containing interventions information registered during 12 years (2006-2017) was provided by the departmental fire and rescue SDIS25, located in Doubs-France.

The analysis of the results demonstrated a high increment in the number of interventions over the years, wherein 12 years this value was more than duplicated. This could represent more and more work for the next years if the pattern is kept. In other words, a change in the management of the budget must be considered to prevent personnel and equipment shortages, to continue improving response times to incidents and to better attend victims' needs. Furthermore, results demonstrated that forecasting firemen interventions with good accuracy are possible and feasible for practical purposes. Considering that both models were basically tuned, better results can be achieved concentrating more efforts on the tuning procedure and on arranging features.

For future work, we will continue testing different Machine Learning methods, combining the LSTM NN with others NN models (e.g., Convolutional NN) and testing a large number of time steps, evaluating and adding new variables to our data set (e.g., social events) and trying out feature selection methods (e.g., F-test and Principal Component Analysis). Additionally, we are working on techniques capable of predicting the sort and the place of interventions to build a complete predictive system for firefighters.

## References

1. France bleu. <https://www.francebleu.fr/infos/societe/60-d-interventions-en-plus-pour-les-pompiers-du-doubs-pendant-la-canicule-1498142401>
2. Ministère de l'écologie, du développement durable et de l'énergie. <http://www.hydro.eaufrance.fr/>
3. Skyfield. <https://github.com/skyfielders/python-skyfield>
4. Alajali, W., Zhou, W., Wen, S.: Traffic flow prediction for road intersection safety. In: 2018 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCom/IOP/SCI). IEEE (oct 2018). <https://doi.org/10.1109/smartworld.2018.00151>
5. Cerna, S., Guyeux, C., Arcolezi, H.H., Lotufo, A.D.P., Couturier, R., Royer, G.: Long short-term memory for predicting firemen interventions. In: 6th International Conference on Control, Decision and Information Technologies (CoDIT 2019). Paris, France (apr 2019). <https://doi.org/10.1109/codit.2019.8820671>
6. Chen, T., Guestrin, C.: XGBoost: A scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 785–794. KDD '16, ACM, New York, NY, USA (2016). <https://doi.org/10.1145/2939672.2939785>
7. Chollet, F., et al.: Keras. <https://keras.io> (2015)
8. Du, S., Li, T., Gong, X., Yang, Y., Horng, S.J.: Traffic flow forecasting based on hybrid deep learning framework. In: 2017 12th International Conference on Intelligent Systems and Knowledge Engineering (ISKE). IEEE (nov 2017). <https://doi.org/10.1109/iske.2017.8258813>
9. Géron, A.: Hands-On Machine Learning with Scikit-Learn and TensorFlow, vol. 1. O'Reilly Media, 1005 Gravenstein Highway North, Sebastopol, CA 95472, 1 edn. (7 2017)
10. Graves, A.: Supervised Sequence Labelling with Recurrent Neural Networks. Springer Berlin Heidelberg (2012). <https://doi.org/10.1007/978-3-642-24797-2>
11. Greff, K., Srivastava, R.K., Koutnik, J., Steunebrink, B.R., Schmidhuber, J.: LSTM: A search space odyssey. IEEE Transactions on Neural Networks and Learning Systems **28**(10), 2222–2232 (oct 2017). <https://doi.org/10.1109/tnnls.2016.2582924>
12. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural Computation **9**(8), 1735–1780 (nov 1997). <https://doi.org/10.1162/neco.1997.9.8.1735>
13. Liu, Y., Wang, Y., Yang, X., Zhang, L.: Short-term travel time prediction by deep learning: A comparison of different LSTM-DNN models. In: 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC). IEEE (oct 2017). <https://doi.org/10.1109/itsc.2017.8317886>
14. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. Journal of Machine Learning Research **12**, 2825–2830 (2011)
15. Shi, X., Li, Q., Qi, Y., Huang, T., Li, J.: An accident prediction approach based on XGBoost. In: 2017 12th International Conference on Intelligent Systems and Knowledge Engineering (ISKE). IEEE (nov 2017). <https://doi.org/10.1109/iske.2017.8258806>