



HAL
open science

Using data science to predict firemen interventions: a case study

Christophe Guyeux, Gaby Bou Tayeh, Abdallah Makhoul, Stéphane Chrétien, Julien Bourgeois, Jacques Bahi

► To cite this version:

Christophe Guyeux, Gaby Bou Tayeh, Abdallah Makhoul, Stéphane Chrétien, Julien Bourgeois, et al.. Using data science to predict firemen interventions: a case study. *Journal of Supercomputing*, 2023, 79 (7), pp.7160-7175. hal-04745432

HAL Id: hal-04745432

<https://hal.science/hal-04745432v1>

Submitted on 20 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Using Data Science to predict firemen interventions: a case study

**Christophe Guyeux, Gaby Bou Tayeh,
Abdallah Makhoul, Stéphane Chrétien,
Julien Bourgeois and Jacques M. Bahi**

the date of receipt and acceptance should be inserted later

Abstract In the recent years, fire departments started to build databases containing detailed information about their interventions during fires, road accidents, and other types of incidents. Their goal is to invest this information using data analysis methods in order to better understand the trends of certain events. This could help them enhance the management of their allocated resources, which leads to a reduction in the operational costs, increase in efficiency and the overall intervention speed. Therefore, in this research paper, we investigate the possibility of predicting future incidents using machine learning algorithms that are trained on a set of data containing information on almost 200,000 interventions that happened during the last 6 years. These data, provided by the fire department in the region of Doubs, France, were not sufficient to detect patterns. Thus, we have imported additional information from external resources that we thought it would improve the accuracy of the predictions. Finally, we tested multiple machine learning algorithms and we compared their results, aiming to determine which algorithm performs better. The results look promising as we were able to predict the number of interventions for each 3 hours block for a whole year, with an acceptable error margin.

Keywords: Machine Learning, Prediction, Data Science, Time series analysis, firemen intervention.

C. Guyeux, G. Bou Tayeh, A. Makhoul, J. Bourgeois, J.M. Bahi are with Femto-St Institute, UMR 6174 CNRS, Université de Bourgogne Franche-Comté, France, email: firstname.lastname@univ-fcomte.fr.

S. Chrétien is with ERIC Laboratory and UFR ASSP, University of Lyon 2, France, email: stephane.chretien@univ-lyon2.fr.

Address(es) of author(s) should be given

1 Introduction

Today's fire-fighters are operating in a technologically progressive environment. Tens of years ago, there was no such thing as a smoke alarm, water sprinkler system, jaws of life, and automatic shut-off valves, etc. While the tools are advancing rapidly, many fire departments are facing budget challenges, rising call volume, personnel and equipment shortages, and the overall expectation to do more with less. Therefore, there is a need for an intelligent system capable of making an assessment of the probability or likelihood that a particular event will occur [1,2]. This can be achieved by building a model that can forecast future events using inputted information about incidents that happened in the past. This intelligent prediction system can help fire departments to manage more efficiently their allocated mobile and personnel resources, enabling them to have the required resources when an incident occur, reduce the response time, and save more lives with less effort [3,4].

Various statistical, data-mining, and machine-learning algorithms are available for use in predictive analysis model in several domains [5–7]. Each of these algorithms was developed to solve specific problems, which may make some of them more appropriate than others depending on the type, size, and other descriptions of the available data. Therefore, it should be run as many algorithms of the appropriate type as possible. Comparing different runs of different algorithms can bring surprising findings about the data. Doing so gives more detailed insight into the problem, and helps identify which variables within the data have the best predictive power.

One of the most known algorithms are the regression ones [8], they can be used to forecast continuous data, such as predicting the trend for a stock movement given its past prices. The linear regression model [9] is one example, it attempts to model the relationship between two or more variables by fitting a linear equation to the observed data. One variable is considered to be an explanatory variable, and the other is considered to be a dependent variable. For example, a modeler might want to relate the weights of individuals to their heights using a linear regression model. Therefore, if we know a person's weight it is possible to estimate its height.

A decision tree [10] is another approach to predictive analysis that is used for prediction and decision making. They are often chosen for predictive modeling because they are relatively easy to understand and effective. The goal of a decision tree is to split a set of data into smaller subsets that are related to each other. Starting at the root which includes the total set of data, and as we move down the tree, the goal is to split them into smaller and smaller subsets at each node of the tree. Each subset must be as distinct as possible from the other in terms of the target indicator. For instance, if we have a set containing information about people, an indicator could be sex, where we split the data into two subsets: one containing females only and the other males. Again, each subset can be split into multiple other subsets based on the age indicator, and so on. The optimal way to do that is by iterating through each indicator as it relates to the target indicator and then choosing the indicator that best splits

the data into two smaller nodes. There are two stages to prediction. The first stage is to build the tree, test it, and optimize it using the available data set. In the second stage, the model is finally used to predict an unknown outcome.

An improved version of the decision tree method is the random forest [11]. It is a supervised learning algorithm that builds a forest consisting of an ensemble of decision trees. The random-forest algorithm brings extra randomness into the model when it is growing the trees. Instead of searching for the best indicator while splitting a node, it searches for the best indicator among a random subset of indicators. This process creates a wide diversity, which generally results in a better model.

One cannot work on machine learning without considering the Support Vector Machines (SVMs) [12]. SVMs are based on the concept of decision planes that separates between a set of data having different class memberships. It was first used for classification purposes, then it showed great performances in Support Vector Regression as well (SVR) [13]. SVR aims to minimize a cost function using a kernel, which could be linear, Gaussian, or polynomial depending on data. The kernel determines the similarity between different features, and thus assign weights to their corresponding cost functions. Features that are close to each other and have the same output will be grouped together due to more weight, while outliers having less weight associated with them are discarded when the cost function is minimized. Thus, outliers will contribute very little to the final predictive model.

Last but not least, the least absolute shrinkage and selection operator (LASSO) [14] is also widely used for prediction analyses. This algorithm works on the concept of penalized regression which helps to select the variables that minimize the prediction error. Ordinary Least Squares regression chooses the beta coefficients that minimize the residual sum of squares (RSS), which is the difference between the observed data and the estimated ones. LASSO adds a penalty to the RSS equal to the sum of the absolute values of the non-intercept beta coefficients multiplied by the parameter λ that slows or accelerates the penalty. E.g., if λ is less than 1, it slows the penalty and if it is above 1 it accelerates the penalty.

In this work, we compare several machine learning algorithms aiming to find the most efficient one in terms of prediction accuracy that estimates the number of interventions for each block of 3 hours. The rest of the paper is organized as follow. Section 2 explains the procedure followed to collect, structure and clean the data. In Section 3 the data are visualized to understand better the hidden patterns and identify and extract the most important features. In Section 4 the previously mentioned machine learning algorithms are tested and their results are compared. In Section 5, a conclusion is provided, with a suggestion to improve the prediction results that opens the doors for a future work.

2 Data Acquisition and Cleaning

2.1 Data Acquisition

The fire department in the region of Doubs-France has provided us with a set of data containing information on a total number of $\approx 200,000$ interventions that occurred during six years, from 2012 to 2017 (included). The data are separated into three different csv files: the list of departures by agents, the list of interventions, and the list of victims. These data files contain information about each incident, such as: The number of intervention (ID), The location, Y and X coordinates, date of intervention, used vehicle and its registration number, departure motivation, alert reception time, departure time, end of intervention time, the total intervention time, age, sex and the state of the victim, etc.

Weather conditions are considered to be a factor that affects significantly the number of road accidents, fires, and casualties. Therefore, including meteorological information to the analysis of incidents trend can improve the prediction results. Moreover, using the previously described csv files, we can extract for each individual intervention the hour of the day when it happened, as well as the day of the week, the month, and the year. This can help us detect tendencies correlated with these parameters (e.g., the number of car accidents increases on Saturday night because young people tend to drink during this period of time). Other parameters that could affect the number of road accidents, fires, and other events can be also taken into consideration, such as traffic hours, academic vacations, holidays, dawn and dusk time, moonrise, moonset, and the phase of the moon as well. The used list of features is provided in Table 1.

The idea is to predict the number of interventions that will occur during each 3 hours time block (since the imported weather data are given by 3h blocks). Therefore, we need to create a dictionary where we can aggregate all the available data provided by the fire department in addition to supplementary data that can be imported from other various sources. In order to build such a dictionary we followed the following steps:

- We initialized a dictionary containing keys ranging from '01/01/2012' until '31/12/17' of the form 'YYYYMMJJhhmmss'. The keys are generated by blocks of 3 hours.
- We imported the following weather-related data from three weather stations located in Dijon-Longvic, Bâle-Mulhouse, and Nancy-Ochey [15]: temperature, pressure, pressure variation each 3 hours, barometric trend type, total cloudiness, humidity, dew point, precipitation last hour, precipitation last three hours, average wind speed for every 10 minutes, bursts over a period, measurement of the burst period, horizontal visibility, and finally the current time.
- We added the imported meteorological information to the previously initialized dictionary. However, the data were not complete, some were missing and marked as 'mq'. Therefore, we applied a linear interpolation to fill the

feature	type	unit
temperature	weather	degree Celsius (float)
pressure	weather	mmBar (float)
wind speed	weather	m/s (float)
wind direction	weather	degree (float)
humidity	weather	percent (int)
dew point	weather	degree Celsius (float)
last hour rainfall	weather	mm (float)
last 3 hours rainfall	weather	mm (float)
horizontal visibility	weather	m (int)
hour	calendar	int
day in the week	calendar	int
month	calendar	int
year	calendar	int
academic vacations	calendar	boolean
holidays	calendar	boolean
dawn time	ephemeris	boolean
dusk time	ephemeris	boolean
moonrise	ephemeris	boolean
moonset	ephemeris	boolean
phase of the moon	ephemeris	float
traffic hours	traffic	int

Table 1 Features used for predictions

- blanks. We then introduce various temporal information such as the day in the week (Monday, etc.), month, year, hour in the day, etc.
- We have extracted the number of interventions from the csv files sent by the fire brigade department. In the latter, there is one line per intervention, which includes the time of the intervention to the second. We group these interventions by a package of 3 hours.
 - The keys 'holidays' and 'startendVacation' are added to the dictionary, which is initialized to 0 (false). The first will increase to 1 for any 3 hours block within an academic holiday period, while the second will increase to 1 for the days corresponding to the beginning and end of holiday periods.
 - We have added the public holidays (1 or 0, for true or false), as well as a second key that is set to 1 the days before public holidays, for the hours ranging from 3:00 pm to 11:00 pm (otherwise 0).
 - We have included information related to the "Bison Futé" which is a system put in place in France to communicate to motorists all the recommendations of public authorities regarding traffic, traffic jams, bad weather, accidents, advice, etc. It classifies the days at risk according to several colors: green = all is well, fluid traffic, orange = dense traffic, red = difficult traffic, traffic jams, black = to avoid because of traffic jams and slow traffic. We integrate these information through two additional keys 'bisonFuteDepart' and 'bisonFuteRetour'. They are 0, 1, 2 or 3, depending on whether the traffic forecasts correspond to Green, Orange, Red or Black.
 - Finally, we added to the dictionary the sunrises, moon phases, etc. A boolean variable 'night' is added to the dictionary to know if it is a day (0) or night (1), for each given hour h. Moreover, we add, for each hour h,

if the moon has risen at h+30min (a Boolean), and what is its phase (an integer from 0 to 7, namely 0 for new moon, 2 for the first quarter, 4 for the full moon, and 6 for last quarter).

In the Table 2 we give an illustrated example showing how the final dictionary looks like. It contains all the information extracted from the provided csv data files (hour, day, etc.) and the ones imported from external sources (meteorological, ephemeris, traffic, vacations, etc.). Each column represents a block of 3 hours. Over the period of 6 years, for each day we have 8 columns representing the 24 hours of the day.'

Table 2 Illustrated example of the final Dictionary

Intervention nbr	0	1	52559
year	2012	2012	..	2012	..	2017
..
StartEnd Vaca- tion	0	0	..	0	..	1
..
WindDirDijon	190	175	..	120	..	100
..
HumBâle	89	85	..	79	..	78
Day	150	150	..	365	..	365
..
3 Hours block	0	3	..	24	..	24
Holidays	1	0	..	0	..	0
VisibiltyNancy	55000	56666	..	48333	..	60000
..
NbrIntervention	7	10		8		6

2.2 Data Cleaning

In this subsection, we will explain how we detected and removed outliers that can affect negatively the end results. At first, we have calculated the mean value of the number of interventions per hour, on average we have 3.59 interventions/hour. Moreover, we found that the minimum number of intervention per hour is 0, the maximum is 85. Finally, in 75% of the cases the number of interventions is less than 5.

Looking at Figure 1, there seems to be some very particular situations, having generated a large number of interventions. As the latter can affect the learning phase, it is appropriate to look at them in more detail, asking whether or not they should be discarded (outliers). We sorted the IDs ranging from 0 to 52559 in descending order according to their corresponding number of intervention. We have noticed that among the top 8 IDs we have 7 neighboring ones (same day, neighboring hours). So we investigated more what happened during this period of time.

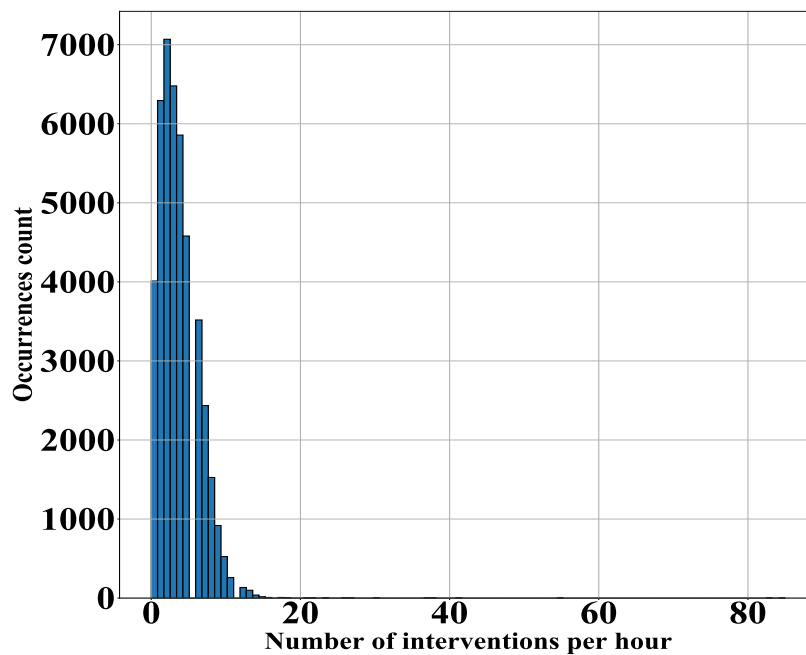


Fig. 1 Histogram showing the frequency of each number of intervention

The ID number 39243 has the maximum number of interventions of 85 interventions. We listed the year, month, day, and the hour of its neighboring ID's. We noticed that they all belong to the night of 24th to the 25th of June 2016. In the csv files, the following main causes were noted for these days: exhaustion, floods, protection of miscellaneous property, and accidents. That particular night there were very violent storms [16], leading to the recognition of the state of natural disaster in the region of Doubs. Therefore, we have two options, either we consider them as outliers and dismiss them in learning by smoothing out the missing data or consider that this is a consequence of exceptional weather, but that with the weather data we should be able to predict this. It remains to be seen whether it seems possible to predict this peak of intervention using meteorological data from Basel, Dijon, and Nancy. In what follows, we will look at an interval of 200 hours (a little less than 9 days) centered around this storm.

We start by looking at the precipitation in the last 1 hour and 3 hours. Figure 2 shows that there is indeed a peak in rainfall during the last hour, but it does not seem obvious: a little less than 4mm, when the article from the 'Est republicain' [16] speaks of 80mm in less than an hour.




Fig. 2 Precipitation each 1h in mm




Fig. 3 Precipitation each 3h in mm

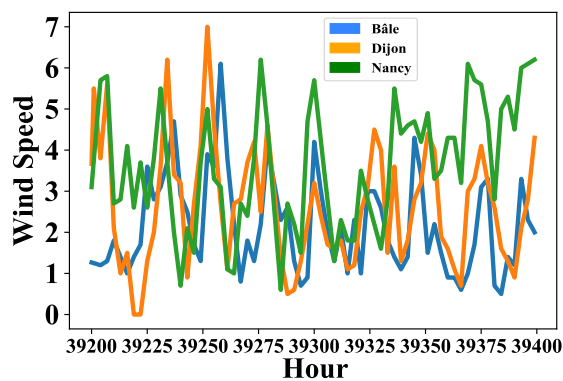


Fig. 4 10-minute average wind speed (m/s)

If we compare the IDs of the peaks with the IDs of the maximum precipitations in the data provided by the weather station in Basel, we notice that they are not among the most important values recorded in this station. This is probably due to Basel's distance from the storm, located between Sancey and L'Isle-sur-le-Doubs (although Basel is closer to this region than Dijon or Nancy). It may, however, be that, at this time, a lot of water has fallen for a relatively long time, so we can look at what it is over a longer period.

Therefore, let us look at Figure 3 at the rainfall over 3 hours. A thunderstorm peak appears clearly, and the amount dropped is closer to the 80mm mentioned above. We checked if such a quantity (a little more than 20mm) is something frequent. It turned out to be the fifth highest rainfall in 3 hours recorded from 2012 to 2016 inclusive. We looked at how many interventions there were in the vicinity of the periods of heavier rainfall over 3 hours. On average, there are more interventions during maximum rainfall over 3 hours. But we remain very far from the number of interventions during the 6 hours of this stormy peak that is considered as a natural disaster.

To conclude on precipitation, one caused an extreme peak in interventions, but this is not the case for the other severe weather events. The most important of these leads to a number of interventions out of the ordinary, but far from the extreme situation studied. These precipitation data are therefore important for our prediction, but they do not allow us to predict the extreme situation of June 25, 2016 (due to weather measurements that are not sufficiently localized).

We look to see if other weather information, measured in Basel, Dijon or Nancy, were remarkable at midnight on 25/06/16. We start with wind speed. As we can see in Figure 4 There is a small peak, but nothing exceptional. The wind speed (10-minute average) was less than 3.5 m/s. We are far from the maximum of 15.9m/s that we can find when sorting the wind speed in ascending order. We also looked at humidity, temperature, and pressure values; there was indeed a drop in pressure, but on the other variables, there was no particular situation at first sight.

It seems difficult, with the weather data currently considered, to predict such a peak of interventions. The situation is sufficiently exceptional and has a real impact on the data considered. For example, the number of interventions in June may be significantly overestimated. For these reasons, we first choose to artificially smooth the intervention data on this date. On outliers, we will put the same number of interventions as the next day at the same time.

In the above, we have detected the presence of outliers by eye, at the level of peaks in the time series. Of course, such an approach is not possible if the algorithm is running in real time in an operational setting, but this can be replaced by an automatic approach, replacing any value beyond 5 times the mean (or any other multiplicative factor deemed relevant by firefighters) with this extreme value. Other more advanced outlier detection techniques can of course also be considered.

3 Data visualization

After cleaning our data and removing the outliers, it is time to analyze it in order to discover the tendencies correlated with the different parameters included in our dictionary. At first, we calculated the degree of correlation of each parameter with the number of interventions. We discovered the following interesting facts: The strongest positive correlation concerns the hour and the strongest negative concerns the night (there are more interventions when it is a day, whatever the season). Moreover, this is something cyclic, every day at 2 am there is less intervention than at 6 pm. The weather data have an impact as well, although less pronounced. Temperature (positive) and humidity (negative: when it rains, people go out less) are first, wind speed and visibility comes second. The eve of public holidays has some importance. The beginning or end of holidays have a greater impact than being on holiday or not. The fact that the moon is apparent plays a little, but less its phase. The year is

also weakly positively correlated: the number of interventions tends to increase from year to year.




Fig. 5 Number of interventions/month




Fig. 6 Number of interventions/Day

Moreover, temperatures are highly correlated with each other and are strongly but inversely correlated with humidity. In a slightly more surprising way, they are correlated with visibility. The relevance of keeping several variables so highly correlated could be considered. If necessary, data could be reduced, for example by aggregating (average, etc.) those correlated, and see if this facilitates learning and predicting the number of interventions.

The first thing we notice is that the number of interventions increases each year, which is normal due to population growth. If we accumulate the number of interventions by month (Figure 5), we notice that the summer months

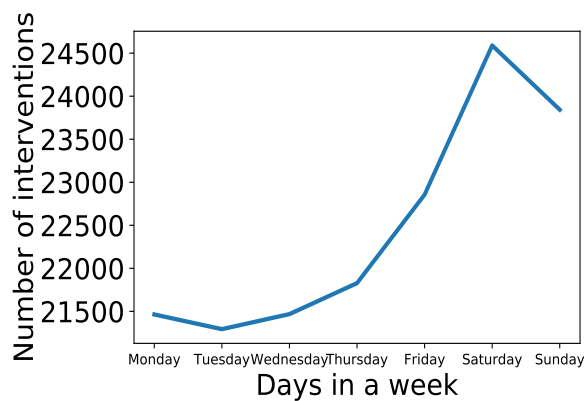


Fig. 7 Number of interventions/Day of the week




Fig. 8 Number of interventions/Hour of the day

are most of the time the busiest in terms of the number of interventions. The end of the years are also busy. One can imagine reasonable causes for this: summer vacations conducive to outings and physical activities, fires. And end-of-year parties. This assumption can be further emphasized by looking at Figure 6, we can clearly see the trends related to summer, as well as a significant number of interventions around mid-February, and lows around April, August-September, and November. Let's move on to the day in the week (Figure 7). Weekends are much busier, with a peak on Saturday (including Friday night after 24h). Let's look at the time in the day (Figure 8). As expected, it is during the hours of the day that the number of interventions is highest, with a small drop during midday. And there is hardly any intervention at 5:00 in the morning.

4 The prediction of interventions

4.1 Learning and testing stages

In this section, we present the best approach that produced the smallest prediction error and we compare the results of the different machine learning methods using the Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) metrics.

In order to perform the prediction, we first split our data into two components, the learning component (2011-2016) and the verification component (2017). The former is used to build a model that will predict the number of interventions for each 3h blocks for the next year (2017), and the latter is used to verify the accuracy of these predictions.

The first step is to specify which data are numeric (year, humidity, day, month, etc.) and which are qualitative (night, holiday, day of the week, etc.), since the latter will be processed by full disjunctive coding when they are normalized (to avoid large data being mixed with small data). An encoder is introduced for qualitative data, for example, if we consider qualitative data at three levels (0, 1, 2), it will be encoded as a three-bit vector, as follows: [1, 0, 0], [0, 1, 0] and [0, 0, 1]. We also create two python pipelines, the first is on the numerical data, allowing them to be normalized (mean 0, standard deviation 1), and the second is on the qualitative data to perform a full disjunctive coding. Finally, the complete pre-processing pipelines for the explanatory variables are formed.

Before testing the different machine learning methods, we start by discovering some reference values concerning the mean square error. Having this information will allow us to have a better estimate of the results produced thereafter. The mean square error, or MSE, reinforces the importance of large errors over small ones and is, therefore, a better score than the mean of absolute errors. We take the square root, leading to the RMSE, to have similar units.

First, we only consider the average number of the interventions for each year, in order to predict the number of interventions for 2017. We obtain the following errors: MAE : 2.39, Root Mean Square Error (RMSE) : 2.94. If we try to do better by taking the average per hour, the following results are obtained: MAE: 1.75, RMSE: 2.28, which is better than the previous results. Let these numbers be our references numbers to test the efficiency of the different prediction models that we will use.

4.2 Prediction

We train the prediction models using the previously described pre-processing pipelines. Afterward, we use the trained model to predict the number of interventions for the year of 2017. We used in this work the Scikit-learn Python module [17] that integrates a wide range of state-of-the-art machine learning

algorithms (including the models that we are interested in) for medium-scale supervised and unsupervised problems. Table 3 summarizes the comparison between the different ML methods, and reveals the best estimator.

Table 3 Error comparison

	RMSE	MAE
Average	2.94	2.39
Average per hour	2.28	1.75
Linear Regression	2.59	1.76
Decision Tree	2.86	2.15
Random Forest	2.19	1.68
SVM	2.21	1.68
LASSO	4.53	3.48

Random Forest: The principle is to drive many decision trees on random subsets of variables and then average their predictions. The comparison between the real number of interventions and the predicted ones is shown in Figure 9. The results look promising since after a random search of best hyper-parameters, the obtained RMSE and MAE are the lowest until now: 2.19 and 1.68 respectively.




Fig. 9 Prediction error using Random Forest

The random forest method achieved the lowest RMSE and MAE. SVM performed good as well, it has the same MAE as the random forest and a slightly bigger RMSE. Surprisingly LASSO achieved the worse results with very high errors. In terms of learning time, the ranking is somewhat reversed. Random forests had the longest learning time, followed closely by SVM. This is not surprising, given the complexity of these methods. LASSO comes next, as an intermediate method. As for the decision tree, its learning time is very low in comparison, which is explained by the fact that random forests are ultimately

only a set of decision trees. The last methods, finally, are immediate. Let us finally remark that a study had been recently conducted for deep learning approaches [18,19]. The conclusion is that deep MLPs or LSTMs do not really improve the prediction quality, but have a much longer learning time. This can be explained by the fact that the information to be extracted from these time series is not huge, and that random forests are a complex enough model to be able to capture all this information.

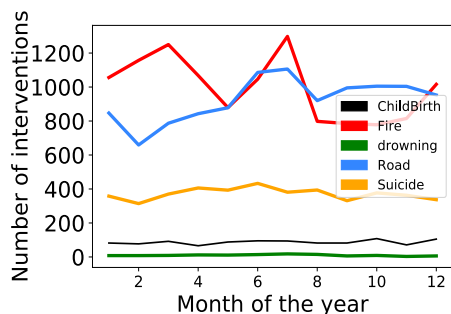


Fig. 10 Number of interventions/month




Fig. 11 Number of interventions/day

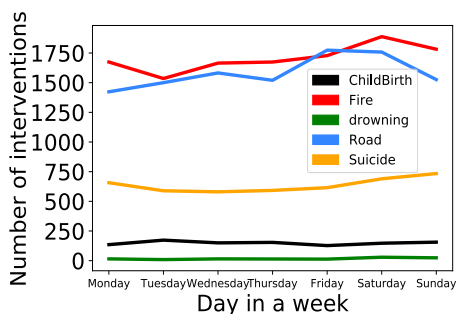


Fig. 12 Number of interventions/day of the week

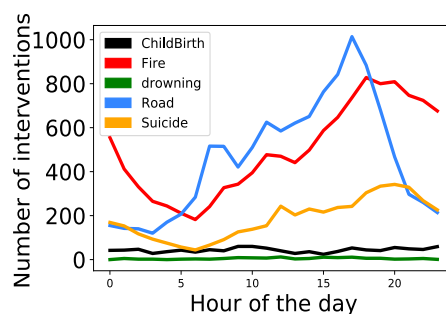


Fig. 13 Number of intervention/hour of the day

5 Discussion and Conclusion

In this work, we are predicting the number of interventions regardless of what is the type of it (suicide, road accident, fire, etc.). The obtained results are fairly acceptable, and we were capable of predicting the number of interventions for each 3 hours block to some extent. However, if we look at Figures 10, 11, 12, and 13, we notice that the number of interventions differs greatly depending on the type. For instance, in Figure 10 we can see a peak in fire interventions starting approximately from May until the beginning of July, and a similar peak for road accidents starts from May until August. Drowning, childbirth,

and suicide remain stable over the whole period. This trend is also visible in Figure 12 where we notice a peak during the weekends in fire and road accident interventions. Moreover, Figure 13 shows a peak between 3 pm and 7 pm for road accidents, and from 4 pm and 8 pm for fire incidents. Therefore, if we can predict first the type of the incident, we will be able to predict more accurately the number of interventions. This indeed is a really interesting approach that we will investigate in more depth in future work.

To sum up this research work, given a large set of data containing information about 200,000 interventions that took place over a 6 years period in the region of Doubs (France), we aimed to determine which technique performs best at predicting the number of interventions for the next 3 hours block with an accuracy that is sufficient for practical purposes. Our results show that the random forest is the best technique for predicting the number of future incidents with respect to the obtained RMSE and MAE values. The predictions are within an acceptable error margin and could help fire departments to anticipate future incidents and better manage their human and mobile resources. The improved management of resources leads to a reduction in the total intervention time and an increase in the response speed. This helps reduce the number of injuries significantly, save more lives, and limit the consequences of an event in the shortest delay possible.

For future work, we will check whether the number of interventions follows a specific probability distribution law and if the associated parameters evolve in time if so we could predict the evolution of these parameters instead of the actual number of interventions. Furthermore, as new data arrive with time, aspects related to incremental learning in evolving domains should be investigated, in order to produce something useful in an operational context. Finally, we will be using Neural Networks for prediction and possibly add additional parameters to the models in order to increase the prediction accuracy.

acknowledgements

This document is the results of the research project supported by the EIPHI Graduate School (contract "ANR-17-EURE-0002"), and the Hubert Curien CEDRE project n° 46543ZD.

Datasets used in this study are available on request from the authors.

The authors have no conflicts of interest to declare.

References

1. Gena R Gerstner, Jacob A Mota, Hayden K Giuliani, Mark A Weaver, Nicholas W Shea, and Eric D Ryan. The impact of repeated bouts of shiftwork on rapid strength and reaction time in career firefighters. *Ergonomics*, pages 1–9, 2022.
2. Wentao Jiang, Chang Chen, and Yan Cai. The design on intelligent physical warning system for fireman. In *2016 International Conference on Audio, Language and Image Processing (ICALIP)*, pages 694–698. IEEE, 2016.

3. N Jayapandian. Cloud enabled smart firefighting drone using internet of things. In *2019 International Conference on Smart Systems and Inventive Technology (ICSSIT)*, pages 1079–1083. IEEE, 2019.
4. Rammah H Eltom, Enan A Hamood, AbdAlrahman A Mohammed, and Abdallah A Osman. Early warning firefighting system using internet of things. In *2018 International Conference on Computer, Control, Electrical, and Electronics Engineering (ICCCEEE)*, pages 1–7. IEEE, 2018.
5. Zaher Mundher Yaseen, Ravinesh C Deo, Ameer Hilal, Abbas M Abd, Laura Cornejo Bueno, Sancho Salcedo-Sanz, and Moncef L Nehdi. Predicting compressive strength of lightweight foamed concrete using extreme learning machine model. *Advances in Engineering Software*, 115:112–125, 2018.
6. Sajid Umair and Muhammad Majid Sharif. Predicting students grades using artificial neural networks and support vector machine. In *Encyclopedia of Information Science and Technology, Fourth Edition*, pages 5169–5182. IGI Global, 2018.
7. Ali Reza Samanpour, André Ruegenberg, and Robin Ahlers. The future of machine learning and predictive analytics. In *Digital Marketplaces Unleashed*, pages 297–309. Springer, 2018.
8. Colin Lewis-Beck and Michael Lewis-Beck. *Applied regression: An introduction*, volume 22. Sage publications, 2015.
9. Douglas C Montgomery, Elizabeth A Peck, and G Geoffrey Vining. *Introduction to linear regression analysis*, volume 821. John Wiley & Sons, 2012.
10. Leo Breiman. *Classification and regression trees*. Routledge, 2017.
11. Andy Liaw and Matthew Wiener. Classification and regression by randomforest. *R news*, 2(3):18–22, 2002.
12. Lipo Wang. *Support vector machines: theory and applications*, volume 177. Springer Science & Business Media, 2005.
13. Alex J Smola and Bernhard Schölkopf. A tutorial on support vector regression. *Statistics and computing*, 14(3):199–222, 2004.
14. Robert Tibshirani, Martin Wainwright, and Trevor Hastie. *Statistical learning with sparsity: the lasso and generalizations*. Chapman and Hall/CRC, 2015.
15. <http://www.meteofrance.com/accueil>.
16. Orages : des inondations aussi dans le Territoire de Belfort. <https://www.estrepublicain.fr/>, 2016.
17. Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011.
18. Selene Cerna, Christophe Guyeux, and David Laiymani. The usefulness of nlp techniques for predicting firefighting responses. *Neural Computing and Applications*, 2021.
19. Selene Leya Cerna Nahuis, Christophe Guyeux, Héber Arcolezi, Raphaël Couturier, Guillaume Royer, and Anna Diva Lotufo. A comparison of lstm and xgboost for predicting firemen interventions. In Reis L. Costanzo S. Orovic I. Moreira F. Rocha Á., Adeli H., editor, *8th World Conference on Information Systems and Technologies*, volume 1160, pages 424–434. Springer (AISC series), April 2020.