Study and predictions of emergency rescue in the Doubs department, France

Christophe Guyeux

August 29, 2022

Abstract

Saving a few minutes in the rescue of a person having an infactus or for a drowning can save lives. The means to save this time are therefore actively sought by the teams in charge of emergency rescue, such as the fire department in France. Part of the answer lies in the fact that some of these accidents have a predictive character: people swim outdoors especially in summer, and when it is hot and the air is polluted, there is an increased risk of discomfort and respiratory distress. The aim of this article is to describe how to experimentally implement an operational predictive tool for emergency rescue. We will study various models, and present the results obtained by the best of them, whose scores effectively allow the solution to be considered operational.

1 Introduction

In a number of countries such as France, the fire department is not only in charge of extinguishing fires, but also has to participate, together with the private ambulance service, in pre-emergency transport. This consists in going as quickly as possible to the scene of the accident, taking care of the patient and, if necessary, taking him/her to the nearest hospital's emergency service. A distinction is made between urgent personal assistance (cardiac arrest, suicide, etc.) and non-urgent assistance (helping an elderly person after a fall, etc.)

Emergency first aid is obviously one of the most important things to consider, since the victim's future is often at stake. Everything that can be done to save time between the beginning of the alert and the arrival on the scene must be investigated. Indeed, a few minutes of delay can cause the victim to fall into a serious condition which, in addition to the additional risks for the victim himself, will result in heavier and longer care, more expensive for society. A few seconds to a few minutes can be fatal in the case of a stroke or a cardiac arrest. This is why the fire and rescue services pay particular attention to optimizing the operational response in the case of urgent personal assistance. One of the most promising avenues is to try to predict the occurrence of such events. For example, predicting that a geographical area is likely to experience such events in an unusual way could lead to the pre-deployment of an ambulance to the scene.

It is clearly unlikely, if not impossible, to predict that an individual will have a cardiac arrest, fall or commit suicide on such and such a day, at such and such a time and in such a place. This does not mean that nothing can be done. For example, during a heat wave, pollution near bus stations and boulevards congested by cars leads to a real increased risk of respiratory distress in sensitive people. Thus, in such periods, if a bus station is such that many elderly people live in its vicinity, then pre-positioning an ambulance in the vicinity does not seem to be a bad idea. And this risk can be predicted, as in modern cities one has access to weather data, air quality data, etc. Other examples can be cited. For example, it has been known since Durkheim [4] that suicides do not happen randomly, but that there are periods of the day, week and year that are more likely to occur. Prepositioning resources seems difficult here, but other actions can be taken, such as imposing agents with resuscitation skills on duty during periods when the risk of suicide is higher than average.

However, artificial intelligence algorithms are now mature, computer hardware has become more powerful, and the computerization of the professional world leads to the existence of learning bases. For example, the French departmental fire and rescue centers have been recording all their interventions (start and end dates and times, type and location of intervention) for several years, for statistical purposes and for legal reasons. Such knowledge bases are beginning to have a sufficiently large history to be successfully exploited when learning a predictive model for emergency rescue.

The aim of this article is to look at a case study, namely the emergency rescue in the department of Doubs, thanks to the exploitation of the data of the departmental fire and rescue service (SDIS 25) to which we had access. The challenge is to explain from start to finish how to set up, in an operational way, a predictive tool that continuously retrieves the variables it needs, and makes the ad hoc forecasts. The model is chosen among the best algorithms of the moment, and the various steps necessary to optimize its results are explained. A discussion is then conducted on the error obtained, and on what can be done at the operational level given such an error.

The remainder of this article is organized as follows.

2 About the data

2.1 The target

We have obtained the data of emergency rescue in the Doubs from the SDIS25, for the period 2015-2021. Once aggregated by hour, we are dealing with a time series of 66968 records (from 01/01/2015 to 22/08/22), whose basic statistics are :

- mean: 1.098465,
- standard deviation: 1.140684,
- minimum 0,
- first quartile: 0,
- median: 1,
- third quartile: 2,
- maximum: 10.

The most frequent slots are those without intervention, when the maximum of 10 was reached only once. And in the vast majority of cases, there are between 0 and 2 interventions, cf. Table 1. An example of the shape of the time series over a few consecutive days is given in Figure 1. The annual trend is increasing, with a slope for the least squares line of 0.0000077, see Fig. 2.

These interventions have systematically increased from year to year, except for 2016 and 2019. But the series, from year to year, follow each other without looking the same. It can be said, however, that from one year to the next, the months of February, April and August are often a little calmer, while December, January, March and September are often a little more hectic (but this is not much more than 20%). Unsurprisingly, there is a peak at Christmas and New Year. Between the least agitated day and the most agitated day, the number of days can be doubled. Some days are less calm than others, and there are obviously local trends, which the sliding average of the Figures 2 and 3 allows to see better. We find that the beginning of the school year (September-October) and the end of the year are agitated, and there are lows during the February and April vacations. A rather calm summer.

One element we look for in time series is the possible presence of an auto-regressive character, i.e. when there is a dependence between what is observed at time t, and what happened at time t-1, t-2... That is to say that the signal is in some way self-sustaining, that local trends exist, which helps in the prediction. Such a study, not illustrated here, shows the effect of local trends. For example:

- if there are no interventions at the present time, then there will be no more than 5 interventions in the next hour,
- when the number of interventions increases at time t, the same applies at time t+1,
- if there were 7 interventions at time t, then there will probably be 4 to 6 interventions at time t+1.

This can be seen through the autocorrelation graph (Figure 4): calculation of the correlation between y(t) and y(t-1) (resp. y(t-2)...). We can see that the current observations are strongly correlated to what happened 1h, 2h or 3h ago (beginning

Nb of interventions	0	1	2	3	4	5	6	7	8	9	10
Nb of hours	24618	22347	12264	5263	1728	550	157	36	4	0	1

Table 1: Number of hours of 0, 1, 2... interventions

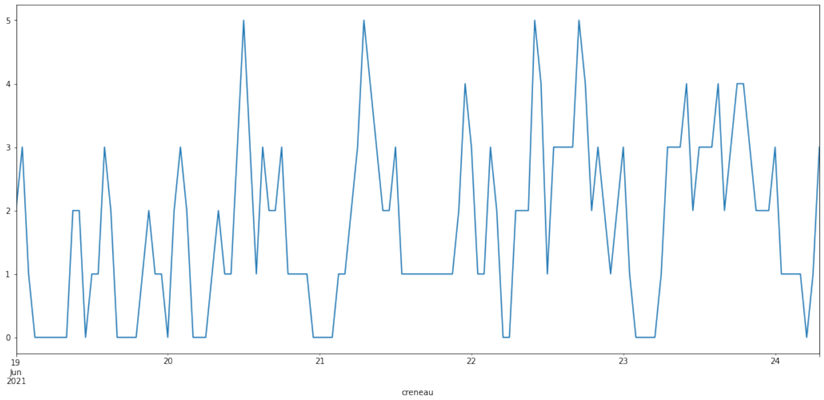

Figure 1: The time series of 6 consecutive days

of the graph). First of all because the interventions initiated at time t-1 have a good chance of not being finished yet at time t, and because there are periods more intense (full day) than others (full night). We also see that at 12 hours of lag, the correlations are negative, although weak: the profile of interventions at 2 a.m. is diametrically opposed to that at 2 p.m. Finally, we can see that there is a daily seasonality: we have more or less the same thing at this time as we had one day ago at the same time.

We therefore have global (annual) and local trends resulting from various seasonality, which can be taken advantage of through calendar variables (year, month, day of the week, day of the year, time) and history (nb of interventions 1h, 2h, 24h ago...)

2.2 The feature

A number of variables are potentially interesting to use as explanatory variables to predict the number of emergency responses.

In our study of the properties of the time series in the previous section, we first saw that there is a fairly pronounced daily seasonality, arguing for the addition of time of day as an explanatory variable. We also noticed a slightly different behavior between working days and weekends, leading us to consider the day in the week as an explanatory variable. The tendency to increase from one year to the next leads to the addition of the year as an explanatory variable. Finally, even if they do not appear clearly, the seasons obviously have an impact that must be considered: we slip on icy patches in winter and drown in a private swimming pool during the summer, and not the other way around. To model this, we will consider the month in the year as an additional feature. It can also be interesting to consider the day in the year, to distinguish for example the 23rd of December from the 24th which, although close in the calendar, have very different intervention profiles. This can be done either by introducing 365 variables, or by means of a target encoding. Finally, the strong autocorrelation of the signal leads us to consider the addition of historical variables (number of interventions one hour ago, two hours ago...)

These features can be deduced directly from the study of the time series. But they are not the only relevant ones. Falling people, floods, road accidents... are clearly impacted by the meteorology: strong winds, heavy rains, snow and ice... Similarly, high temperatures lead to discomfort and dehydration. So it seems sensible to add such features.

Other features seem relevant. Concerning floods, the heights and flows of the rivers in the region are also interesting to consider. For road accidents, road traffic information is interesting, if it is avail-

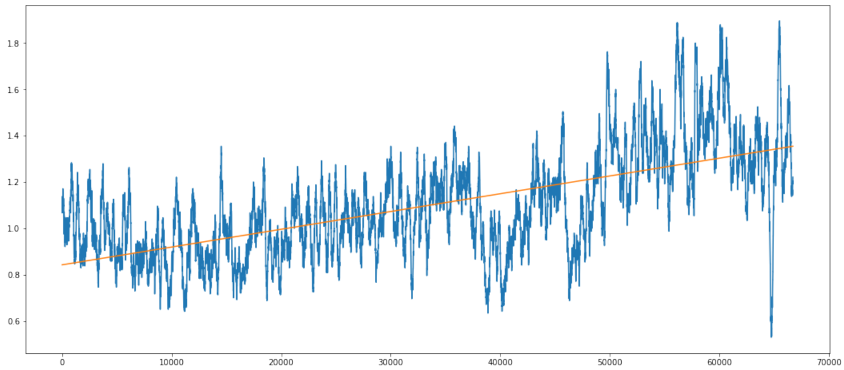

Figure 2: Shifted rolling mean (window=240) of the time series

able (with history and in real time). For ailments, air quality plays a role, and is frequently available online, at least for cities of some importance. School vacations should also be taken into account. Finally, it can be interesting to monitor the evolution of local searches on google, such as the keywords flu, diarrhea, or gastroenteritis for the prediction of epidemics which, if they come back every year, do not come back at the same date.

3 Implementation of the predictive tool

Among the a priori relevant variables, it is first necessary to sort out the features that are really exploitable for a predictive model, and those that, although interesting, cannot be exploited. Such features are only useful if, on the one hand, we have access to a sufficiently large history, and if, on the other hand, we have easy and continuous access to new values of these variables. The first condition is necessary to be able to do learning, while the second is sine qua non to do online, continuous prediction. For a given feature, these two conditions can be satisfied in some countries, or even in some cities, and not in others, which leads to a customized deployment in each case.

For our part, we had access to each of the features considered in the previous section, both historically and continuously, with the exception of road traffic data. Pipelines had to be set up both for the constitution of the knowledge base and for the continuous collection of variables for online prediction. Such a pipeline produces, in our case, a dataframe with 1299 columns, after standardization of numerical features and encoding of qualitative ones. However, some of the variables in this dataset are highly redundant or correlated with each other, while others have almost no predictive power. For example, the amount of water that has fallen in the last 3 hours and the height of the rivers are redundant. Similarly, the temperature is somehow contained in the month in the year. And providing a large number of variables, with high correlations, does not help the model to do quality learning.

Feature selection is an important step in the process. Various techniques exist, and none is perfect. Each one follows a well-defined approach, and only sheds light on the correlations from a certain angle. This is why we have opted for a consensus approach, where we have kept only the features that seem to be important for the majority of the selection approaches, the approaches being: linear correlation coefficient, variance threshold, univariate variable selection test returning F-statistic and p-values [8]... The result of such a selection is an ordered list of features according to their importance, consisting here of the number of approaches for which they are important.

In our previous studies [6, 7, 5, 2] of the best artificial intelligence algorithms for firefighting prediction, we concluded that XGBoost [3] achieved the best compromise between learning speed and prediction quality. We will therefore retain it here, and compare it with naive methods and with the AR approach. Similarly, we have seen that the "maximum depth" hyperparameter was the most important to

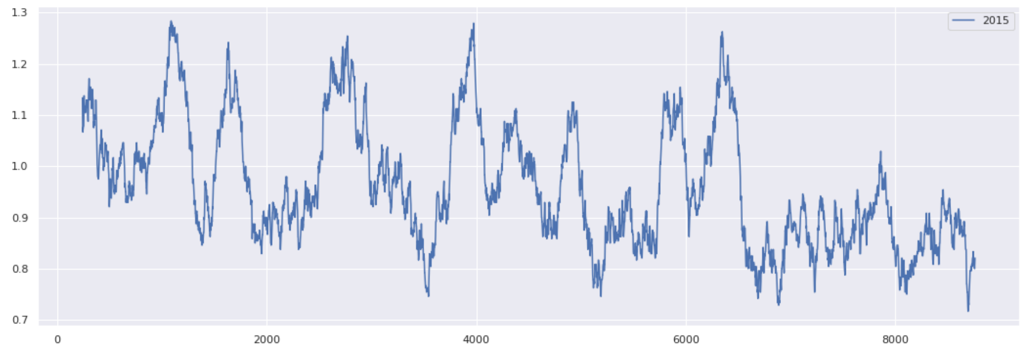

(a) 2018

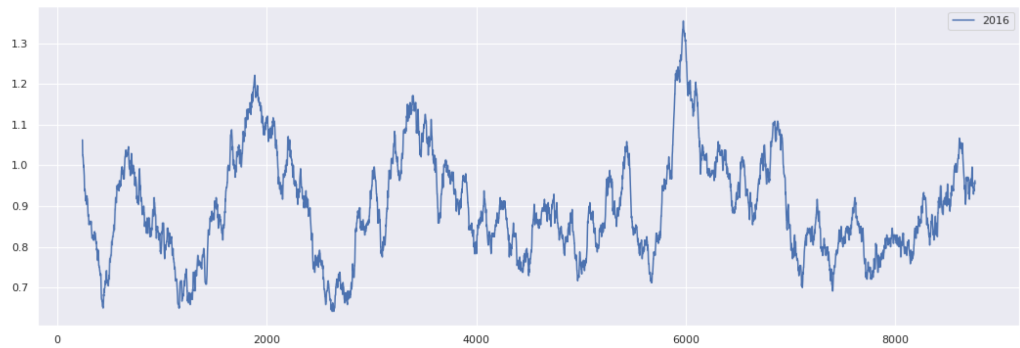

(b) 2020

Figure 3: Comparison between two years of interventions (rolling mean)

vary, the other hyperparameters having a lesser and redundant influence in the case we are interested in. We then proceeded to a gluttonous approach at the level of the features classified according to their importance.

We start from the first variable of the list, and we give it the best Root Mean Squared Error (RMSE) score obtained when we vary max_depth from 2 to 15. At iterate k, we have retained k features, and we test the addition of a k+1-th feature by testing in order the remaining features, and by varying each time the max_depth as above. We stop the process when it is no longer possible to improve the RMSE by adding a feature to those selected. Once the set of features is definitively adopted, we proceed to a refinement of the score by playing on various hyperparameters of XGBoost, using the Optuna tool [1].

4 Obtained results

Let's start by getting some baseline values, which will allow us to measure how much better we can do than these naive approaches. Considering the mean as the baseline predictor, we get an RMSE of 1.14. And if, instead of always predicting the mean 1.05, we prefer to predict the mean per hour to take into account the daily seasonality, finding a more or less sinusoidal signal around the mean, we find an RMSE of 1.09. As for the persistence model, duplicating the last known value, it allows us to do better, because of the auto-regressive character of the signal. It leads to an RMSE of 1.05. The signal is better, but we are still late, which is very penalizing for important changes in two successive hours. Finally, we can mix the two previous models (average per hour and persistence), to try to take advantage of both the seasonality and the auto-correlation. The predicted signal computed as

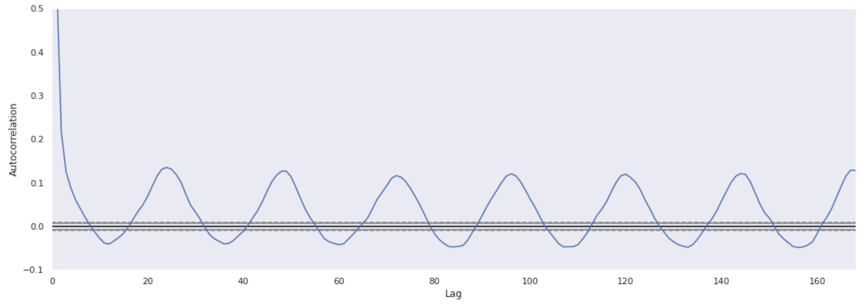

Figure 4: Autocorrelation graph of the time series

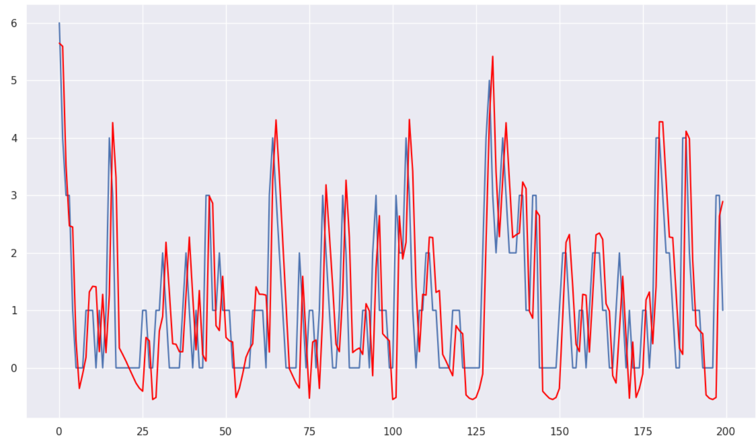

Figure 5: Mixed approach between average per hour and persistence

the average of the two previous models looks more interesting (cf. Figure 5), but the score is not as good (RMSE of 1.097), because we do not do as well on the stagnation periods.

We can try to have a pure time series approach, looking only at the auto-correlation of the signal (cf. Figure 4), but in a finer way using the AR model [9] (correlation of 0.576215 between t and t+1 values). An autoregressive AR-X(p) model order selection leads to a max lag of 50. With such a lag, the obtained RMSE is 0.922, which is not as good as the best approaches seen so far. But the shape of the predicted signal is not so bad, see Figure 6.

This brings us to the model that is currently in production. We have chosen 23 explanatory variables, including features of history, weather, google searches, as well as rolling mean, number of interventions in progress in the last hour, current time, day of the week and day of the year. We have set the hyperparameters of XGBoost as detailed above. We considered a Poisson type regression, since it is a counting process. In doing so, we obtained an RMSE of 0.764, much better than what we had obtained until now. As can be seen in Figure 7, such an RMSE is low enough to make the predictions operationally useful.

5 Conclusion

In this article, the possible predictability of emergency rescue has been studied. The objective was to know if, as a whole (natural hazards, accidents, respiratory distress...), this type of intervention possessed elements of recurrence, tendency or seasonality sufficiently marked to make them, up to a certain point, predictable, and if the use of external features could help in such predictions. The series was studied in depth, which made it possi-

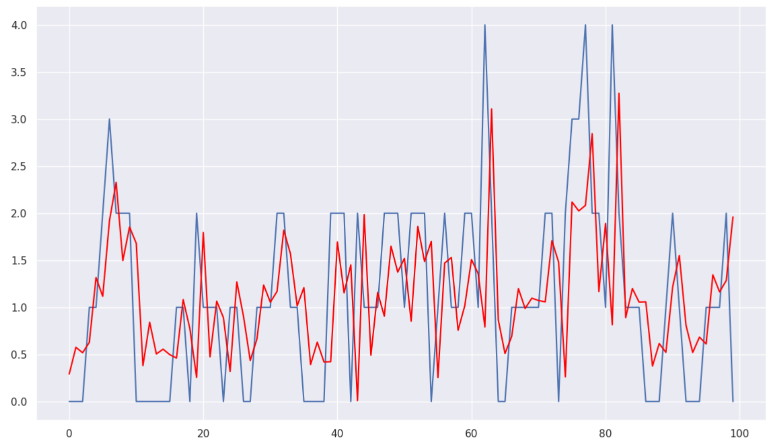

Figure 6: Predictions by using an AR model

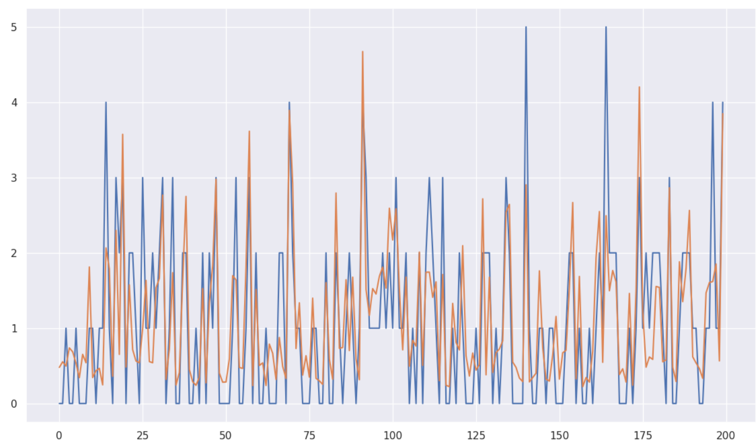

Figure 7: Predictions by using XGBoost

ble to highlight a certain number of features to be exploited. The features of interest were described. The study was carried out up to the constitution of the predictive model, its training, and its evaluation. The results show that such a prediction is feasible, and that the scores are good enough to have an operational interest.

Our next work consists first of all in confirming these results by using other data sets. We will investigate other cities, and other time horizons, to see if this approach can be generalized. We will also be interested in non-emergency rescue, which, although less important than emergency rescue, is still of some interest. Each natural hazard will be studied separately, to see if ad hoc features could not improve the quality of predictions, and we will also be interested in industrial hazards. The final goal will be to produce a professional quality software, integrating all these aspects, and useful from an operational point of view.

References

[1] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.

[2] Selene Cerna, Héber Arcolezi, Christophe Guyeux, Guillaume Royer-Fey, and Céline Chevallier. Machine learning-based forecasting of firemen ambulances' turnaround time in hospitals, considering the covid-19 impact. *Applied Soft Computing Journal*, 109, September 2021.

[3] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.

[4] Emile Durkheim. Suicide: A study in sociology (ja spaulding & g. simpson, trans.). *Glencoe, IL: Free Press.(Original work published 1897)*, 1951.

[5] Christophe Guyeux and Jacques M. Bahi. How to predict patient arrival in the emergency room. In *10th World Conference on Information Systems and Technologies*, pages ***–***, April 2022.

[6] Christophe Guyeux, Abdallah Makhoul, and Jacques Bahi. How to build an optimal and operational knowledge base to predict firefighters' interventions. In *Intelligent Systems Conference*, pages ***–***, September 2022.

[7] Roxane Mallouhy, Christophe Guyeux, Chady Abou Jaoude, and Abdallah Makhoul. Machine learning for predicting firefighters' interventions per type of mission. In *8th International Conference on Control, Decision and Information Technologies*, pages ***–***, May 2022.

[8] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[9] Robert H Shumway, David S Stoffer, and David S Stoffer. *Time series analysis and its applications*, volume 3. Springer, 2000.