

PREDICTIONS IN PRE-HOSPITAL EMERGENCY TRANSPORT IN FRANCE : A STATE OF THE ART

Christophe Guyeux

ABSTRACT

For a number of years now, the regional fire department centers have been recording their interventions numerically. Such databases are under-utilized and are mainly used for statistical and management purposes. However, such a history of interventions can be very useful, if used in conjunction with artificial intelligence algorithms, for predictive purposes. Such work has recently been done in France through a series of articles investigating the various aspects of the problem, and has been put into production at the Doubs center. The objective of this review is to take stock of all the work that has been done so far, to list the successes and the stumbling blocks, and to draw up a roadmap on this theme for the years to come.

1. INTRODUCTION

Depending on the country, and in addition to their firefighting missions, firefighters can be assigned various tasks related to personal assistance. In France, for example, they are one of the essential components of the emergency rescue service, and are called upon for missions as diverse as removing bodies from crashed vehicles, lifting elderly people who have fallen, or for respiratory distress. They share these missions with private ambulances and some ambulances available in hospitals (called SMUR).

These fire departments have been facing an increasing number of requests in the last few years, which they are finding more and more difficult to meet. This is due to several causes such as the ageing of the population, the disengagement of the private sector in actions considered as not very profitable, the increase of pollution and global warming at the origin of many respiratory problems, etc. On the other hand, the human and material means devoted to the fire departments do not increase any more, for reasons of rationalization of the public expenditure. In this context, the optimization of the use of existing means seems to be one of the most promising approaches to allow the fire departments to continue to fulfill their mission without too much degradation of the quality of service.

A promising way to achieve such optimizations is to use load prediction, which will reduce the duration of interventions, thus relieving the pressure on the demand for resources. This reduction can be achieved, for example, by pre-positioning ambulances in high-risk areas, thus reducing the time to arrive at the scene. It also allows for the redistribution of resources to available centers according to expected demand.

However, such a prediction assumes that learning has taken place, and therefore that a learning database exists. In order to build such a database, one must first have a sufficiently large history of past interventions, of sufficiently good quality. Such a history exists within the Service Départemental d'Incendie et de Secours du Doubs (SDIS 25), and various researches and proofs of concept could be realized from this data set. The objective of this paper is to show what has been possible to achieve from this history, proving that useful predictions can be made provided that

such a dataset is available. In what follows, we will first present in detail the SDIS25 intervention history, and its various aspects of trend, seasonality, autocorrelation... allowing us to hope that machine learning is possible and that it leads to good predictions. We will then review the various issues that have arisen and the various

Achievements that have been published in the last decade. The presentation of these results will be followed by a discussion, in which we will mention the various avenues for future improvement.

2. DATA ANALYSIS

Most of the departmental fire and rescue services have a digital device allowing them to save all their interventions: start and end time, type and location of the intervention. For privacy reasons, these interventions are aggregated by type and time, and only the hourly streams leave the centers.

A basic descriptive statistical analysis leads to the following results for the Yvelines department:

- an average of 32.4 interventions per hour, a standard deviation of 18.2,
- a minimum of 0 interventions, a maximum of 224,
- the median is 32: 50% of the hours have less than 32 interventions,
- the first quartile is 18: 25% of the slots have less than 18 interventions,
- the third quartile is 45: only 25% of the hours have a number of interventions greater than or equal to 45.

When we study these flows as time series, we see that there is a daily seasonality, see Figure 1. This is well justified by the fact that each day there is more activity, and therefore more accidents, during the day than at night; a decrease in activity during the lunch break; and more activity at the end of the day than at the beginning. These trends are found every day of the week, and we see from one week to another, a certain difference between workdays and weekends (slight shift of activity to the evening). These trends alone show that prediction of interventions is possible. Finally, there is no clear weekly seasonality, see Figure 2.

Similarly, the relationship between months in the year and overall level of interventions is not clear. In fact, it is true that a number of types of response occur only at a limited time of the year, such as influenza or gastroenteritis episodes, hymenoptera outbreaks, or ice storm accidents. But these episodes do not occur on the same days (or even weeks) from one year to the next, and the accumulation of the various types leads to an averaging of interventions that drowns out the underlying signal, see Figure 3. However, taken separately, each of these events has a certain predictive character.

3. INVESTIGATIONS CARRIED OUT

3.1. Regression for Sdis25 Interventions

Being able to predict the daily activity of firefighters is of great interest to optimize human and material resources. This will allow a faster response by achieving a better geographical deployment of these resources according to the number of planned interventions. After obtaining the list of interventions for the period 2012-2017 in the department of Doubs, France, we added a relevant collection of explanatory variables based on calendar data (time of day, day of week, day of month, year, holidays, etc.), road traffic, weather and astronomical data, etc. [10]. After detecting outliers and filling in missing data, this set was divided for training, validation and testing. The training is then performed on an ad hoc multi-layer perceptron (a basic neural network) whose

hyperparameters have been finely defined using our computational mesocenter. This neural architecture is finally applied on a real case, namely the prediction of firefighters interventions for the year 2017 after a learning phase on 2012- 2016, leading to really encouraging first results.

The above study was repeated, but on the consolidated database of interventions over the period 2015-2020 [13]. The objective of this work was then to test various predictive models usually studied in the literature, and to see if the results obtained above by the neural network could be beaten by methods such as the simple decision tree, random forests, the ADA Boost algorithm, vector support machines, and Bayesian approaches. At the end of this study, neural network and decision tree ensemble approaches were found to be the best.

Following the previous study, we decided to further investigate the neural network approach [12]. Recurrent networks are well suited to the data considered, since the number of interventions at time

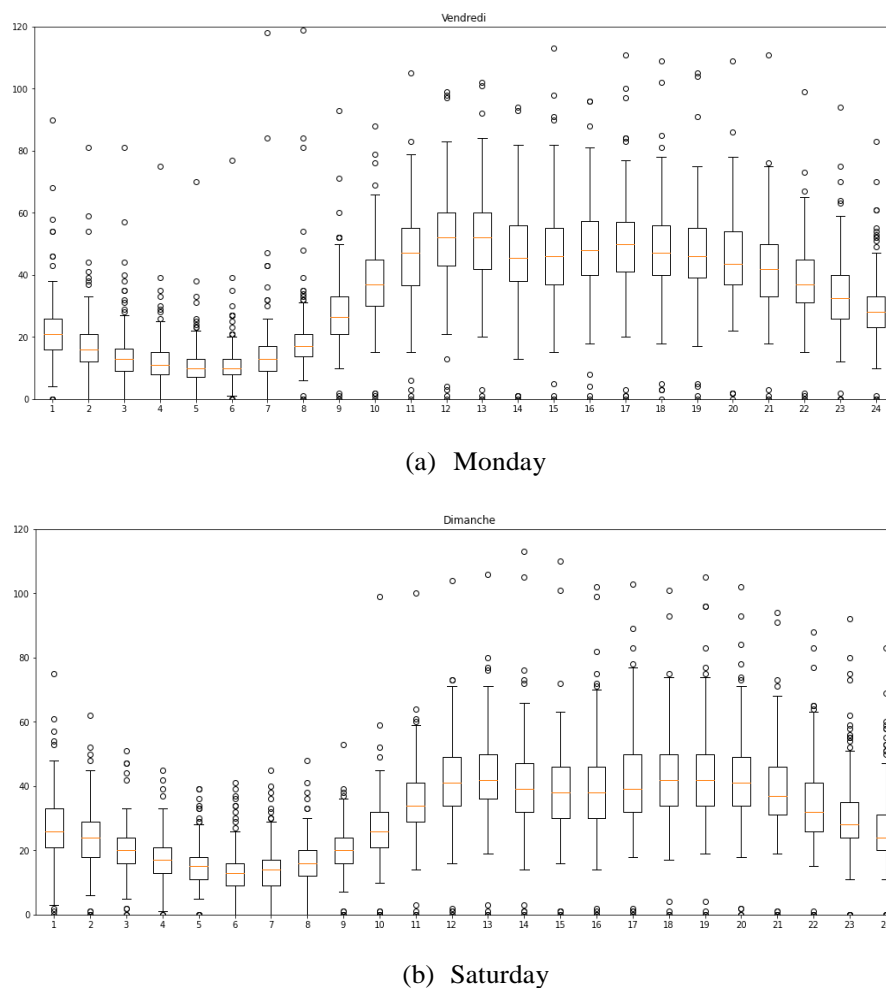


Figure 1. Daily intervention flow profile (number of interventions per hour

depends, to a large extent, on the number of interventions at time $t - 1$. Long Short-Term Memory (LSTM) has recently been shown to be the champion among recurrent neural networks, notably because it avoids the disappearance of gradients. In this paper, we have shown that these LSTMs provide better prediction results than all the models tested so far.

The paper [11] aims to show that machine learning tools are finally mature enough at the present time to allow useful predictions, even when considering rare events such as natural disasters. We have once again considered LSTM, which we have compared to XG Boost, a state of the art method based on a set of decision trees. The whole process is detailed, from data collection to predictions. The results obtained prove that such a quality prediction is quite feasible by both methods, even to a certain extent in case of rare events such as natural disasters, but that it could still be improved, for example by optimizing the hyperparameters of these algorithms. Finally, we note that the learning process for XG Boost is faster than for LSTM.




Figure 2. Seasonal decomposition of the intervention time series

While human-related interventions are quite predictable, the recognition of extreme events such as storms or floods requires, as seen above, more information to reach a high level of reliability. Therefore, [5] proposes to improve the above by applying automatic natural language processing (NLP) techniques for extracting texts from weather reports and processing them to identify rare events. In a first step, the previous regression models were improved by taking into account new explanatory variables, namely the alert level of weather reports (color code: green, orange, red). In a second step, using a multi-label classification approach, we took advantage of the texts of the bulletins themselves, again with LSTM, which we compared to convolutional neural networks (CNN), and then to very recent transformers made for NLP such as Flaubert and Camembert, to predict the peak periods of interventions for 4 categories. The results in recognizing intervention peaks due to rare events were improved when texts were taken into account. The best model obtained was the Camembert model, with an accuracy of 92% for heating and 86% for storm/flood, compared to 71.44% and 63.77% for the previous approaches.


The article [17] is a natural continuation of the studies on rare events, extending the analysis to all types of events. Fire department operations vary according to weather, climate, season, etc. For example, the frequency of traffic accidents is greater during the day than at night. Adjusting the needs to the demand by category of operation would allow firefighters to have a head start by knowing the type and number of human resources needed for the next interventions. The objective of this study is to see to what extent the number of firefighters' interventions can be predicted for 12 different categories, namely: childbirth, drowning, wasps, brawls, street fire, suicide, flooding, traffic accident, heating, fire, aid to persons, and emergency aid to persons.

Pre-emergency transport, whether performed by firefighters, SMUR or private ambulances, has a predictive character due to the fact that rescue is directly related to human activity, which is itself predictable. XG Boost has emerged as the best tool to predict the number of interventions by type, but how to design an optimal and operational knowledge base has not been discussed until now. In [9], we explain how to realize such a base with a content that is both relevant and can be continuously updated, making possible the industrialization of the process, and thus a better


operational response of the concerned services. We show that three feature selection tools custom-built for XG Boost are mature enough to allow the optimization of such a database, and a good accuracy in predictions. We also show what these tools can bring in terms of business knowledge, and discuss the organizational and efficiency consequences that such an optimized predictive model could bring.

3.2. Concerning the Arrival at the Hospital

Most hospital emergency departments are suffering, faced with a constant increase in the number of patients without the financial means to follow them.



(a) 2018



(b) 2020

Figure 3. Comparison between two years of interventions

Part of the answer to this worrying situation lies in the optimization of existing resources, for which artificial intelligence techniques show promise. In [8], we evaluate this possibility in a concrete way, starting from real data and applying a comparative analysis of 4 state-of-the-art algorithms. An original method of selecting the explanatory variables is applied, and the hyper parameters of the algorithms are also precisely chosen. The results show that the most powerful machine learning algorithms currently available are quite capable of making good predictions of the number of patients arriving at the emergency room, provided they are well applied. When ambulance turnaround time (TT) in emergency departments is prolonged, it not only affects the victim severely, but also leads to unavailability of resources in emergency medical services (EMS). This problem can be aggravated with abnormal situations, such as the current coronavirus 2019 (COVID-19) pandemic. Taking this into consideration, [4] presents a first study on the impact of COVID-19 on ambulance TT by analyzing SDIS 25 historical data for three hospitals. This article also

calculated and analyzed the number of breakdowns in the services, which increased due to the lack of available ambulances again.

Therefore, it is vital to have a decision support tool to better reallocate resources by knowing the utilization time of ambulances and pre-emergency services personnel. Thus, this paper proposes a novel two-step methodology based on machine learning (ML) models to predict the TT of each ambulance in a given time and hospital. The first step uses a multivariate regularly spaced time series model to predict the average transport time (AvTT) per hour, which takes into account temporal and external variables (e.g., COVID-19 statistics, weather data...). The second step uses a multivariate irregularly spaced time series model, which takes into account the time variables of each ambulance departure, the type of response, the external variables and the previously predicted AvTT as inputs. Four state-of-the-art ML models were considered in this paper, namely Light Gradient Boosted Machine (LightGBM), Multilayer Perceptron, LSTM and Facebook Prophet. The proposed methodology provided usable results in practice. The AvTT accuracies obtained for the three hospitals were 90.16%, 97.02% and 93.09%. The accuracy of TT was 74.42%, 86.63%, and 76.67%, with a margin of error of ± 10 minutes.

3.3. Predictions on Anonymized data

Predicting the number and type of interventions of civil security services is essential, both to optimize the size and competence of firefighters on call, and to pre-position material and human resources... To accomplish this task, it is necessary to have artificial intelligence skills, which are not usually found in a medium-sized fire department. Such a request can be mandated, for example, to specialized companies or research laboratories. This mandate requires the transmission of potentially sensitive information related to interventions, victims... which are not intended to be made public. The objective of [6] is to show that a machine learning tool can be deployed and provide accurate results, using a learning process based on anonymized data. Learning on real but anonymized data was performed using extreme gradient boosting (XG Boost) based on the latest studies, and the performance of each anonymization was compared on the number and type of interventions per day.

The work [2] focuses on predicting the number of firefighter responses in certain locations while respecting the strong concept of Differential Privacy DP. A local differential privacy approach was first used to anonymize the location data. Statistical estimators were then applied to reconstruct a synthetic dataset uncorrelated to users. Finally, a supervised learning approach using XG Boost was used to make the predictions. Experiments showed that the anonymization and prediction method is very accurate: introducing noise to sanitize the data does not affect the quality of the predictions, and the predictions accurately reflect what happened in reality.

Emergency medical services (EMS) provide critical emergency assistance and ambulatory services. One of the key measures of EMS service quality is ambulance response time (ART), which is generally the time between EMS notification and when an ambulance arrives on the scene. Because many victims require care within an adequate time frame (e.g., cardiac arrest), improving response time is vital. [3] proposes to predict ARTs using machine learning techniques, which could be used as a decision support system by EMS to enable dynamic selection of ambulance dispatch centers. A well-known predictor of ARTs is the location of the emergency (e.g., whether it is an urban or rural area), which is a sensitive input because it can reveal who received care and for what reason. Therefore, we considered the "input perturbation" setting in the privacy-preserving ML literature, which allows the EMS to sanitize each location data independently and, therefore, the ML models are trained only with anonymized data. In this paper, geo-indistinguishability was applied to sanitize each emergency location data, which is a state-of-the-art formal notion based on differential privacy. To validate our proposals, we used SDIS25 retrospective data as well

as public data (e.g., weather and traffic data). As the results show, sanitizing location data and perturbing their associated features (e.g., city, distance) did not significantly impact ART prediction. Therefore, one can share sanitized datasets to avoid potential data leakage, membership inference attacks, or data reconstruction, for example.

[1] proposes a new methodology based on machine learning (ML) techniques to predict both the mortality of victims and their need for transport to health facilities from data collected from the beginning of the emergency call to the SDIS25 notification. We first analyzed SDIS25 calls to identify associations between call handling times and victim mortality and to measure the importance of the variables. Next, we validated our ML-based methodology, where mortality could be predicted with an accuracy and AUC of 96.44% and 96.04%, respectively, while need for transportation scored 73.62% and 78.91%, respectively. Furthermore, we found that it was still possible to predict both values by perturbing the input data, applying k-anonymity and differential privacy techniques.

4. TIME SERIES APPROACH

We have seen that managing explanatory variables on a continuous basis is a complex matter, so we can question the relevance of their use. In other words, the sequence of interventions is a time series that has its own trend, seasonality, self-sustainability, etc., and perhaps a prediction based on these elements alone can lead, in a much simpler way, to something that is still usable in practice. [14] aims at studying three different time series prediction algorithms and comparing them to the predictions of previous works. The results obtained are interesting, leading to the questioning of the relevance of a coupling (average, etc.) of the two approaches: regression with explanatory variables, on the one hand, and time series (without explanatory variables) on the other hand.

[7] focuses on the issue of the prediction horizon. In the above, the intervention database was enriched with explanatory variables conditioning human activity (and thus the occurrence of accidents, and then of firefighters' interventions): meteorological, epidemiological, calendar and ephemeris data, holidays.... In the last models, there were 1600 such variables, and a selection of models led to the choice of the XG Boost, or the LSTM if the learning time is not decisive, the latter being able to learn the number of interventions for a given hour, knowing the values of the explanatory variables mentioned above. This approach has a major flaw, making it unusable in an operational way for firefighters, because we need to know the values of the explanatory variables (meteorological, epidemiological, etc.) of a given hour to predict the number of interventions. Therefore, only past hours, for which the values of the variables are known, are available to predict the number of interventions. The objective of this work is to explain how to adapt this so that we can make predictions of firefighter responses in the future, at a time when the explanatory variables are not available. We start by changing the target value of "number of interventions per hour corresponding to the explanatory variables" to "number of interventions per hour +1", "per hour +2", etc. and study the evolution of the error with the prediction horizon. We then implement techniques based on the history of the explanatory variables, such as duplicating for the coming year the values of the variables of the past year, or the average of the previous years, to see the evolution of the prediction error compared to the previous method. Some variables can be accurately calculated in the future, such as the daily time of dawn and dusk or the phases of the moon, and thus the prediction of future interventions can be done by restricting to these exact variables. Finally, we will attempt to predict first the evolution of the explanatory variables, and then that of the target variable (the number of interventions at the chosen horizon), considering each explanatory variable as a time series, and using traditional techniques for predicting such series (e.g., AR, MA, and ARIMA) or more recent ones like Prophet. The best coupling between the methods of constitution of the future explanatory variables and the machine learning techniques will thus be determined and compared to a pure time series analysis of the interventions (as above), in

order to arrive at an efficient aid solution compatible with the distance of the prediction horizon.

In addition to the AR, MA, ARIMA and Prophet methods discussed above, there are time series prediction methods based on "exponential smoothing", which may be interesting to study in view of the seasonality of the intervention series. Experiments have shown [16] that the Holt-Winters method has the best accuracy compared to the classical time series prediction techniques previously studied and to other exponential smoothing methods.

In the last work based on time series analysis techniques [15], we looked at the extent to which anomaly detection and breakpoint detection techniques can be of practical interest for the operational management of crises such as those of COVID. The objective is twofold. Firstly, to see if we can signal a change in trend or behavior in the series of interventions in order to detect the onset of a crisis with some time in advance. Secondly, to improve the quality of learning by eliminating anomalies and by breaking down the time series into single elementary trends (and multiplying the number of learnings accordingly).

5. DISCUSSION

While the study of the predictability of firefighting operations has progressed well over the past five years, a certain amount of work remains to be done in order to reach maturity in this field of research, and the widespread use of such a predictive tool.

First of all, if the prediction of calls, or of any type of intervention, at the level of the whole department gives satisfactory results, this is not the case for sub-types of intervention such as emergency rescue in big cities, or chimney fires. Predictions for such subtypes can be very useful, however. For example, there are so-called "hot" neighborhoods in the two large cities of the Doubs, which for social reasons are more prone to urban violence. This violence leads to a large number of accidents, fires, etc. in a short period of time and in a very limited area, and being able to anticipate such situations would make it possible to size the guards, or even to pre-position human and material resources. It could be possible to predict this risk to a certain extent, as these incidents frequently have an origin (police blunder, defeat at a soccer match, etc.). The means of such a prediction could be natural language processing on news from the local press and on the agitation at the level of social networks, but until now everything remains to be done.

At this level, the key is to find use cases, to find situations where predictions can be both accurate and useful. For example, as we have seen in the various articles presented here, predicting any type of intervention at the geographical level of the department and for the hour to come is done very well, and the error is really minimal. But such a prediction has no practical use. On the other hand, if we change the target from "any type of intervention" to "any phone call received by the fire department", we remain on something easily predictable, but which gains in utility. Indeed, it allows to size the call room staff, and to have more people to pick up the phone in peak periods. The usefulness must therefore guide the choices of what to predict, in consultation with the people in the business. But it cannot be the sole driver of the research, and must be coupled with feasibility. For example, predicting the risk of chimney fires is useful and feasible at the level of the department or of fairly large regions, but will never be feasible at the level of a village. Similarly, predicting the risk of illness at the street level would certainly be very useful, but can only be envisaged if we have explanatory variables that are both very precise and very localized, and with a significant history, and we risk waiting decades before having this.

One of the first works that can be carried out combining both a real usefulness and a certain feasibility concerns crisis situations associated with natural or exceptional risks. We have carried out a first work in this sense at the level of the risk of flooding in the Doubs. This risk comes from

natural rivers which, due to particular climatic conditions, have a level and flow that increase. Once a certain level is reached, the rivers overflow and the fire department faces a peak of interventions over a short period of time. These peaks are directly linked to local weather data, which is available in real time. Even better, there is a whole network of flow and height measurements of various rivers, which are also available in real time. It is obvious that these flood intervention peaks are directly and strongly correlated to these explanatory variables, and the operational interest of such predictions is also obvious. These predictions are not without difficulty, however, because they are rare events and because we have only a small history. But these predictions can only improve over time (there is another difficulty coming from the fact that the climate is changing, in these times of global warming).

Various other hazards, natural and otherwise, are also useful to predict. Examples include forest fires, agricultural fires, coastal storms, avalanche risk, and even industrial risks. These risks are not the same from one region to another, and each has its own dynamics. It is therefore necessary to find the right explanatory variables each time and to select an ad hoc model. However, some risks are much more difficult to predict than others. For example, the floods in the Doubs mentioned above are easy to predict, because there is a real inertia, a real slowness in this event: we can observe three days in advance, sufficiently upstream of the river (the Doubs), the flood to come in Besançon. Conversely, the so-called Cevennes events taking place in the mountains in the south of France lead to spectacular rises in mountain rivers in an instant, and are therefore very difficult to predict, although dangerous.

Predictions of such risks are therefore useful, but only if a certain accuracy is achieved. To do so, we would need a lot of data, especially since we are dealing with rare phenomena. However, this is the main pitfall in this problem: the departmental fire and rescue services have only recently digitized their intervention data, and we only have at best a 7-year history for the most advanced departments. For flood interventions in the Doubs, this translates into about ten events, which is clearly insufficient to achieve good predictions on various time horizons. On the other hand, the fact that the risks are not the same from one department to another prevents, to a certain extent, to compensate for the low temporal depth of the history by a geographical multiplicity from the hundred departments of France. Unless one is willing to wait for the history to grow in time, the solution is to group the data from centers with similar profiles, after a clustering step. However, this solution has yet to be implemented. However, there is a final obstacle at this level, which for the moment has no solution: the world is evolving and the dynamics of interventions are changing over time, for various reasons related to global warming, aging of the population, new epidemics such as Covid-19, etc. Also, the history does not fully reflect what is happening now.

Once solutions to these problems are found, it would be useful to broaden the focus, and extend the predictions to all pre-emergency transportation. In France, the latter is shared by three entities: the fire department, private ambulances, and public ambulances of hospital emergency services. Each of them has an impact on the others. Thus, strikes in emergency departments, the closure of small hospitals, or the reduction of licenses for the private sector, lead to an overload for the departmental fire and rescue services, and taking into account the one makes it possible to better predict what will happen to the others. And taking into account the pre-sanitary sector as a whole makes it possible to have a longer-term visibility, and for example to distribute the right number of approvals for the private sector, so that the latter can relieve the public sector while being economically viable (a fair balance must be found). Finally, emergency pre-hospital transport is not only of interest in France, and the same problems are found, *mutatis mutandis*, in other countries. We should be able to move on to the international level, knowing how to extract what is common to each country, and what must be adapted locally.

6. CONCLUSION

In this article, we have summarized our work on the prediction of firefighting interventions in France. A complete review of the literature has been carried out, including both explanatory variables based approaches and those focusing on time series. This was followed by a discussion of current successes and areas for improvement. The latter will be the subject of future publications.

In this future work, we will first focus on the prediction of natural hazards, which have a strong operational impact. We will also see how to aggregate data from various departmental fire and rescue services, in order to increase the knowledge base in cases where the type of intervention to be predicted is a rare event. Finally, we will try to extend the study to the whole of the pre-sanitary emergency transport, and for other countries than France.

REFERENCES

- [1] Heber H. Arcolezi, Selena Cerna, Jean François Couchot, Christophe Guyeux, and Abdallah Makhoul. Privacy-preserving prediction of victim's mortality and their need for transportation to health facilities. *IEEE Transactions on Industrial Informatics*, 2021.
- [2] Heber H. Arcolezi, Jean-François Couchot, Selene Cerna, Christophe Guyeux, Guillaume Royer, Béchara Al Bouna, and Xiaokui Xiao. Forecasting the number of firefighters interventions per region with local-differential-privacy-based data. *Computers & Security*, 96, September 2020.
- [3] Héber Arcolezi, Selene Cerna, Christophe Guyeux, and Jean-François Couchot. Preserving geolocalizability of the emergency scene to predict ambulance response time. *Mathematical and Computational Applications*, 26(3), September 2021.
- [4] 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.
- [5] Selene Cerna, Christophe Guyeux, and David Laiymani. The usefulness of nlp techniques for predicting firefighting responses. *Neural Computing and Applications*, 2021.
- [6] Jean-François Couchot, Christophe Guyeux, and Guillaume Royer. Anonymously forecasting the number and nature of firefighting operations. In *Proceedings of the 23rd International Database Applications & Engineering Symposium*, pages 30:1–30:8. ACM, June 2019.
- [7] Christophe Guyeux. Predicting the firemen interventions: a concrete case study. In *European Network for Business and Industrial Statistics*, pages ***–***, September 2019.
- [8] 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.
- [9] 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.
- [10] Christophe Guyeux, Jean-Marc Nicod, Christophe Varnier, Zeina Al Masry, Nouredine Zerhouni, Nabil Omri, and Guillaume Royer. Firemen prediction by using neural networks: a real case study. In Kapoor S. Bi Y., Bhatia R., editor, *Intelligent Systems Conference*, volume 1037, pages 541–552. Springer, Cham, September 2019.
- [11] 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.
- [12] Selene Leya Cerna Nahuis, Christophe Guyeux, Héber Hwang Arcolezi, Anna Diva Plasencia Lotufo, Raphaël Couturier, and Guillaume Royer. Long short-term memory for predicting firemen interventions. In *6th IEEE International Conference on Control, Decision and Information Technologies*, pages
- [13] Selene Leya Cerna Nahuis, Christophe Guyeux, and Guillaume Royer. Artificial intelligence to assist and optimize fire brigade services. In *The International Emergency Management Society, 2020 Annual Conference*, pages ***–***, November 2020.

- [14] Roxane Mallouhy, Christophe Guyeux, Chady Abou Jaoude, and Abdallah Makhoul. Time series forecasting for the number of firefighters interventions. In The 35-th International Conference on Advanced Information Networking and Applications, pages ***-***. Springer Series "Advances in Intelligent Systems and Computing", May 2021.
- [15] Roxane Mallouhy, Christophe Guyeux, Chady Abou Jaoude, and Abdallah Makhoul. Anomalies and breakpoint detection for a dataset of firefighters' operations during the covid-19 period in france. In 10th World Conference on Information Systems and Technologies, pages ***-***, April 2022.
- [16] Roxane Mallouhy, Christophe Guyeux, Chady Abou Jaoude, and Abdallah Makhoul. Forecasting the number of firemen interventions using exponential smoothing methods: a case study. In 36-th International Conference on Advanced Information Networking and Applications, volume 449, pages ***-***. Springer LNNS, April 2022.
- [17] 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.