



The usefulness of NLP techniques for predicting peaks in firefighter interventions due to rare events

Selene Cerna¹ · Christophe Guyeux¹ · David Laiymani¹

Received: 30 March 2021 / Accepted: 30 January 2022 / Published online: 26 February 2022
© The Author(s), under exclusive licence to Springer-Verlag London Ltd., part of Springer Nature 2022

Abstract

In some countries such as France, the number of operations assisted by firefighters has shown an almost linear increase over the years, contrary to their resource capacity. For this reason, predicting the number of interventions has become a necessity. Initially, time series models were developed with several types of qualitative and quantitative features, including the alert level of the bulletins, to predict the operational load. We realized that interventions related to human activities are quite predictable. However, the recognition of interventions due to rare events such as storms or floods needs more than quantitative meteorological data to be identified, since there are almost always zero cases. Thus, this work proposes the application of natural language processing techniques, namely long short-term memory, convolutional neural networks, FlauBERT, and CamemBERT to extract features from the texts of weather bulletins in order to recognize periods with peak interventions, where the intense workload of firefighters is caused by rare events. Four categories identified as Emergency Person Rescue, Total Person Rescue, interventions related to Heating, and Storm/Flood were our targets for the multilabel classification models developed. The results showed a remarkable accuracy of 80%, 86%, 92%, and 86% for Emergency Rescue People, Total Rescue People, Heating, and Storm/Flood, respectively.

Keywords Firemen activity prediction · Regression · Natural language processing · BERT

1 Introduction

The missions and the activity of fire departments change from one country to another: In some countries, fire departments are only in charge of extinguishing fires, while in others they are also in charge of rescuing people, whether it is urgent or not. In countries such as France, fire activities represent only almost 10% of all their interventions, and they may be called out for road accidents, floods, respiratory ailments, gastroenteritis, or even wasp nests. In western countries with an aging population, the share of personal assistance is constantly increasing, when repeated economic crises lead to budget cuts.

Thus, in countries where fire brigades activity encompasses rescuing people, the number of interventions has been steadily increasing for several years now and the COVID-19 pandemic has only amplified the situation. In this context, ensuring rapid and efficient interventions becomes a major challenge for many brigades and it seems interesting to try to plan interventions in advance. Better still, answering the questions When? Where? Of what type (road accident, domestic accident, flood, etc.)? can have a strong impact and therefore save lives. This difficult problem is just beginning to be studied [1–3]. In fact, the vast majority of brigades have accumulated an important mass of data related to their past interventions. It is therefore interesting to try to use these data and the recent advances in machine learning (ML) to try out intervention prediction. It is reasonable to think that such predictions are possible, as the reasons for these interventions are to some extent deterministic. Forest fires occur more frequently in dry, hot weather than in wet, rainy weather; floods follow heavy rains; domestic and work accidents occur mostly in the middle of the day, and they are rare at 3 am; falls on the ice do not occur in summer and drowning

Authors have contributed equally to this research work and David Laiymani is the corresponding author.

✉ David Laiymani
david.laiymani@univ-fcomte.fr

¹ Femto-ST Institute, UMR 6174 CNRS, University Bourgogne Franche-Comté, Belfort, France

in outdoor pools does not occur in winter. As can be seen from the above examples, weather conditions have an undeniable impact on human activity and the accidents it causes, and therefore on the activity of firefighters. In addition, most national weather prediction services provide online services or APIs to retrieve various physical quantities useful for weather knowledge. These quantities, which are internationally codified and are mandatory measurements, include temperature, pressure, wind direction and strength, dew point, or hydrometry, for a set of stations spread over the entire territory.

However, while these physical quantities are useful in predicting, to some extent, the weather, they are only partially effective in predicting firefighters interventions, their types, and intensities. Knowing that a thunderstorm event is coming does not accurately predict the extent to which firefighters activity will be affected. More precisely, by coupling the history of weather conditions with the history of interventions, we can see that while some storm-type weather conditions clearly lead to peaks in intervention, others, on the contrary, go somewhat unnoticed. In other words, the simple quantitative value of the physical quantities to be measured does not alone make it possible to accurately predict the occurrence of periods of heavy intervention (for example, heavy rains that do not systematically lead to flooding). In this article, we would like to insist on the fact that most states have also set up a meteorological vigilance service, with reports indicating the risk incurred (snow, ice, storm, floods, etc.), the duration of the event, and its textual description, as well as useful advice. It is clear that integrating the level of vigilance in the form of qualitative variables, one for each risk monitored, allows on the one hand to improve the regression scores in the learning phase of the total number of interventions. On the other hand, we argue that the textual content of newsletters is valuable and underexploited, that sentences such as “it is advisable to unplug electrical appliances” or “it is strongly advised not to walk along the coast” are rich in information for the prediction of interventions such as heating or emergency rescue due to rare events, and that automatic natural language processing tools are now mature enough to understand the link between such assertions and peak intervention periods.

1.1 Purpose and contributions

With these elements in mind, the present work performs a detailed study of various ML models developed for the prediction of intervention peaks due to rare weather events and identifies the remarkable impact of NLP models and weather bulletin texts. The rarity of events such as flood or heating events makes them difficult to predict due to the small amount of data available and, as we will see in the

remainder, these events can be the source of extremely high numbers of interventions. The predictions are made for four categories of interventions, namely Emergency Person Rescue, Total Person Rescue, interventions related to Heating, and interventions related to Storm/Flood. The data were collected from 2012 to 2020. They are composed of interventions of the Fire Department of Doubs (SDIS 25), located in the northeast of France, quantitative weather variables and texts of weather bulletins and their vigilance levels from Météo-France. In this way, SDIS 25 or other emergency medical services (EMS), in general, could better recognize periods with high workload generated by rare events, through text processing of weather bulletins, and strategically prepare the appropriate personnel and armament to deal with the crisis so that no service disruptions occur and more lives can be saved. In summary, this article proposes four main contributions described as follows.

- a. *Analysis and prediction of interventions using basic univariate time series models* This allows to recognize the seasonality of rescue-type interventions and identify the lack of recognition of incidents triggered by rare events. We compared three basic models per category to predict the number of interventions per hour. Each model makes predictions equal to the mean, equal to the last known value, and equal to the mean per hour.
- b. *Analysis and prediction of interventions using multivariate time series models* This allows to complement and identify the trend and the seasonality of the signal by adding qualitative and quantitative variables (vigilance indicators and weather and calendar variables). However, these are still not sufficient to recognize interventions caused by rare events. We developed seven models per category with the state-of-the-art extreme gradient boosting (XGBoost) technique, in which the features are combined to measure their impact on the prediction.
- c. *Analysis and prediction of intervention peaks using multilabel classification models based on decision trees and tabular data* The problem is restated as a multilabel classification task for the four categories, using the variables of the previous model and the XGBoost and random forest (RF) techniques well known in the literature for their fast execution and robustness. This allows to determine the influence of the variables according to the new perspective and deduce that there is still a lack of information to recognize the peaks due to rare events. Thus, these models become our baselines for the following models with natural language processing (NLP).

d. *Analysis and prediction of intervention peaks using multilabel classification models based on NLP techniques and text from meteorological bulletins* We developed and compared models based on ancient (Long Short-Term Memory – LSTM and Convolutional Neural Network – CNN) and modern (FlauBERT and CamemBERT transformers) NLP techniques and only texts extracted from weather bulletins. This allows to significantly improve the forecast of the peaks compared to the previous models with decision trees and tabular data and thus demonstrate that it is possible to extract much more information from public weather bulletins using NLP techniques.

The structure of this article continues with Sect. 2 that reviews the contributions of related works. Section 3 describes the acquisition and preprocessing of the data. Section 3 presents the types of neural networks applied for the natural language processing of bulletins. Section 5 presents the results of three basic approaches developed before the approach with NLP techniques, to analyze and demonstrate the efficacy of the latter with the best model found. The article ends with Sect. 6, in which our conclusions are shown and avenues for future improvements are discussed.

2 Related work

Among the works reviewed and related to the optimization of fire departments responses to incidents, we mainly found contributions for the prediction of interventions [3, 4] and fires [5]. Likewise, in certain parts of the world firefighters are part of EMS, since they also provide ambulance services. In this way, predictions of traffic accidents [6], ambulance response time, and resource allocation are also included [7–9]. Furthermore, we can find works related to the predictions of rare events such as earthquakes [10] and hurricanes [11], which would allow firefighters to identify a specific location for damage assessment and develop better strategies when succoring the population.

In fact, the aforementioned works did not use NLP techniques. However, in other studies, NLP techniques demonstrate their outstanding utility by enriching data sources and predictions. This is achieved by recovering habits, preferences, emotions, feelings, and distress messages through the recognition of semantic patterns [12–14] from various media such as social networks, the news, therapeutic reminders, information systems under the video-based, audio-based, and text-based formats [15]. On the one hand, in [16], the authors presented a cognitive assistant system for EMSs based on Google Speech API, in

which the voice records (incident description and patient status) received by the respondent are converted into texts for extract medical concepts. In this way, the system responds by providing information to rescuers on protocols to follow such as resuscitation and airway management. On the other hand, NLP has contributed to the generation of terminological sources for the classification and forecasting of rare events or crises [17–20]. For example, in [18], it is presented the first study for crisis management using French transformer-based architectures (BERT, FlauBERT, and CamemBERT) apply to French social media, to classify tweets for natural disasters. In [21], the authors make use of the Bayesian model averaging approach and linear-chain conditional random fields to extract knowledge from tweets and build a decision support system to identify early warning signs of earthquakes. Also, [22] developed the Flood AI Knowledge Engine, which is a system composed of ontology management, query mapping and execution, and NLP modules. The system provides emergency preparedness and response, as well as knowledge about flood-related resources for the population. This knowledge is returned by interpreting natural language queries from users.

Having seen the impact of NLP on the recognition of extreme events and given that, to the authors' knowledge, no previous studies have exploited a source such as weather bulletins to detect trends in the number of firefighter interventions, the present work takes advantage of NLP techniques to process these bulletins and predict peak intervention periods for the categories: Emergency Person Rescue, Total Person Rescue, interventions related to Heating, and Storm/Flood.

3 Data retrieval

3.1 Data acquisition

We collected the interventions of the SDIS25 firefighters over the period 2012–2020, for the entire department, for each time slot of this period (76224 slots, 1 slot = 1 hour), and the following four types of intervention: Emergency Person Rescue, Total Person Rescue, interventions related to Heating, and Storm/Flood. A short statistical description of each type of intervention (per hour) is provided in Table 1, when Fig. 1 plots the curves for each type for the first days of 2012.

It can be seen that in both types of personal assistance (emergency and total), the number of interventions is on average high enough to show seasonal patterns: The 24-hour cycle is clearly shown in Fig. 1. We can also see that some days are special, such as New Year's Day and that the integration of calendar variables into this daily

Table 1 Statistical description of the interventions per hour

Type of intervention	Mean	Std.	Min.	Max
Emergency rescue of people	3.56	2.46	0	20
Total rescue of people	7.25	4.53	0	30
Heating	0.32	0.62	0	6
Storm and flood	0.26	1.17	0	82

seasonality should improve the predictions. Finally, the rarity of storms or heating events will make their predictions problematic in the absence of additional information, which already shows the interest of considering weather-type variables. This is even more true for the “storms and floods” case as it has the lowest average (0.26 interventions per hour) and an extremely high maximum (82 interventions in one hour): These events are very rare, but the source of an extremely high number of interventions.

This is the reason why we have retrieved historical meteorological data from the Météo-France site (essential SYNOP data [23]). The three closest meteorological stations selected are those of Nancy-Ochey (latitude 48.581000, longitude 5.959833), Dijon-Longvic (47.267833; 5.088333), and Basel-Mulhouse (47.614333; 7.510000). The data recovered in this way are temperature (degrees Celsius), pressure (Pa), pressure variation (Pa per hour), barometric trend (categorical), humidity (percent), dew point, last hour rainfall (millimeter), last three hours rainfall (millimeter), mean wind speed (10 min., m/s), mean wind direction (10 min., m/s), gusts over a period (m/s), horizontal visibility (m), and current weather (categorical, 100 possible values).

These data were supplemented by (textual) vigilance alert bulletins from Météo-France [24]. They are XML files containing the type of vigilance (heatwave, extreme cold, snow or ice, thunderstorms, strong winds), the


beginning and end of the vigilance period, the level of the alert (green, orange or red), a detailed description of the risk (including the locations impacted, the conditions to be expected, etc.), as well as a set of very detailed advice to users. An example of such files is provided in Fig. 2. Finally, we added calendar-type variables, namely the time of the niche considered, the day in the week, the day in the month, the month in the year, and the year considered.

3.2 Data preprocessing

Totally, 12218 weather bulletins have been produced since 2011, including 1054 for the northeast region (known as CMIRNE) of France, which interests us. Eighteen departments are represented in this region, but we were only interested in the Doubs and its three neighboring departments (Territory of Belfort, Haute-Saône and Jura) in a first time.

The set of texts available may seem small at first glance, but on the one hand, each bulletin covers a number of departments and has several sections (location, description, qualification of the phenomenon, new facts, current situation, expected evolution, possible consequences, and behavioral advice), and each section is made up of several long and detailed sentences. The result is a corpus of 76333 characters.

These bulletins are then cut out by section. Over the entire period of vigilance, the average number of interventions is calculated for each of the four types considered, and a class 0 or 1 is introduced depending on whether this number is above the average number of interventions of the type in question, for all the period 2012–2020. We thus associate 4 binary labels for each text in each section of each bulletin, as shown in Table 2. Through this encoding, we became interested in the question: “in the context of

Fig. 1 Interventions for the four types considered, early 2012

```

<Titre name="Conséquences_possibles">
<Paragraphe>
<Intitule>Vent/Orange</Intitule>
<Texte>
* Des coupures d'électricité_et_de_téléphone_peuvent_affecter_les_réseaux_de_distribution
pendant_des_durées_relativement_importantes.
</Texte>
<Texte>
* Les_toitures_et_les_cheminées_peuvent_être_endommagées.
</Texte>
<Texte>
* Des_branches_d'arbre risquent de se rompre.- Les véhicules peuvent être déportés.
</Texte>
<Texte>
* La circulation routière peut être perturbée, en particulier sur le réseau secondaire
en zone forestière.
</Texte>
<Texte>
* Quelques dégâts peuvent affecter les réseaux de distribution d'électricité_et_de_téléphone.
</Texte>

```

Fig. 2 Example of a possible consequences section (in French)

this vigilance event, should we expect an increased number of interventions such as heating, or storm, etc.?”. In other words, by this increased number, we mean a higher than the average number of interventions.

In Fig. 3, one can observe in more detail the number of samples (texts) with value 0, where the number of interventions was below or equal to the mean, and value 1, in the opposite case. Furthermore, we find an unequal distribution in each binary class of each category that might bias the prediction models.

4 NLP models

Let us now introduce the neural networks that have been considered here for natural language processing.

4.1 LSTM

Long short-term memory (LSTM) networks are a type of recurrent neural network (RNN) which is a category of neural networks dedicated to sequence processing [25]. In

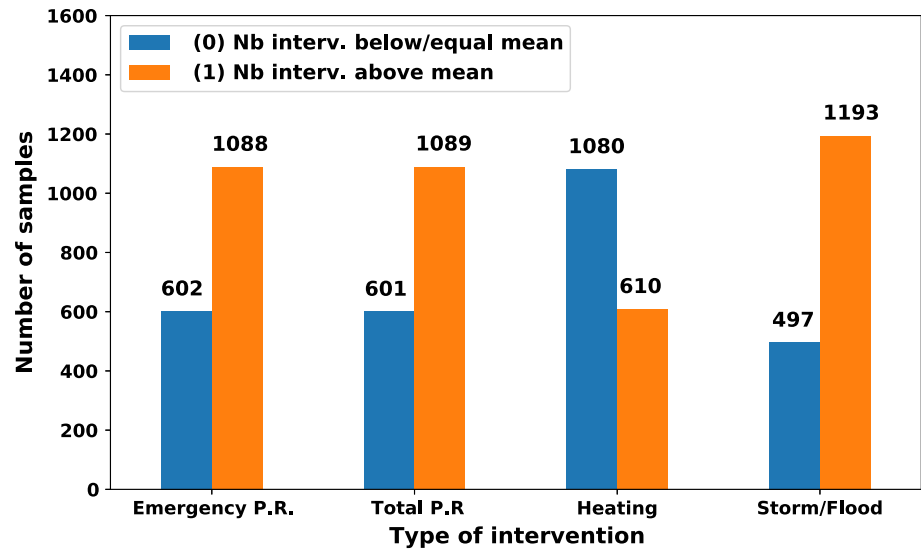
the case of NLP, RNN are interesting since, first they can process sequences of variable size and second the use of recurrent connections allow to analyze the past part of the signal. In this way, RNN are particularly well suited to handle three different types of problems: sequence labeling, sequence classification, and sequence generation. A recurrent network can be approximated by a non-recurrent network unfolded in time. But as the unfolded network is deeper, the vanishing of the gradient is more important during learning, and it becomes more difficult to train. In the same way, as the weights of the recurrent layer are duplicated, RNNs are also subject to exploding gradient. Although they are very effective for modeling short- or medium-term dependencies, RNN are still insufficient for modeling long-term or very long-term dependencies. In NLP, it is common to need to model dependencies of the order of a hundred or more time steps. That is why LSTM has been introduced.

To model very long-term dependencies, it is necessary to give recurrent neural networks the ability to maintain a state over a long period. This is the purpose of LSTM cells. The cell can be seen as an internal memory and can

Table 2 Example of a multilabeling of vigilance texts

Text	Emergency person rescue	Total person rescue	Heating	Storm/flood
Les températures sont déjà négatives aujourd’hui mercredi. A 15h les. une vigilance particulière notamment pour les personnes sensibles ou exposées.	1	0	1	1
Période de grand froid; moins intense qu’en 1985; mais nécessitant toutefois. températures sous abris observées s’échelonnent entre -1 et -4 degrés.	1	0	1	1

Fig. 3 Number of samples for each category



maintain a state for as long as necessary. It consists of a numerical value that the network can control depending on the situation. The memory cell can be controlled by three control gates: The input gate, i_t , decides whether the input should modify the content of the cell; the forget gate, f_t , decides which memories will be eliminated from the previous long-term state c_{t-1} ; and the output gate, o_t , decides whether the cell content should influence the output $y(t)$ of the neuron. Let x_t be the present entry and h_{t-1} the preceding short-term state. LSTM can be expressed in the following way:

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

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

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

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

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

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

where W_{xi} , W_{xf} , W_{xo} and W_{xg} are the weight matrices for their connection to the input vector x_t ; W_{hi} , W_{hf} , W_{ho} and W_{hg} are the weight matrices for their connection to the previous short-term state h_{t-1} ; and b_f , b_g , b_i and b_o are the bias terms of each layer.

LSTM are generally used in layers and in this case, the outputs of all neurons are fed back to the inputs of all neurons.

In this paper, we used the Keras and TensorFlow libraries in Python to initially built several architectures with 1, 2, and 3 LSTM layers, different numbers of neurons, and constant values for learning rate and batch size. Then, we evaluated their performances and selected the three architectures that gave us the first best results, these

are the ones described in Table 3. Finally, to obtain the best LSTM model, we intensified the search with the three selected architectures, varying the learning rate and batch size parameters. For this, we used the HyperOpt library with 100 iterations and the Tree Parzen Estimator suggest (tpe.suggest) algorithm, which models two density functions instead of the probability of an observation to estimate the expected improvement of a new configuration.

Furthermore, before the texts enter the neural network, a preprocessing was applied as detailed below:

- Words with flexible endings were eliminated from the French texts and their base forms were returned, this is known as lemmatization, through the “spacy” library. Unlike a pre-trained model in a large vocabulary, this LSTM model uses the texts of the weather bulletins as a base, it is there the need to bring the words to their base form to identify a greater risk of a possible event. For example, the words “inondations” and “inondables” share the same root and refer to “floods.”

Table 3 Defined architectures for LSTM

Archi. 1	Archi. 2	Archi. 3
Input(200)	Input(200)	Input(200)
Embedding(100)	Embedding(100)	Embedding(200)
LSTM(128)	LSTM(1000)	LSTM(128)
Dense(256, ReLU)	Dense(2000, ReLU)	Dense(256, ReLU)
Dropout(0.5)	Dropout(0.5)	Dropout(0.2)
Dense(4, Sigmoid)	Dense(4, Sigmoid)	LSTM(512)
		Dense(1024, ReLU)
		Dropout(0.2)
		Dense(4, Sigmoid)

- b. Using the “nlk” library, French stopwords were eliminated, that is, words such as articles, pronouns, prepositions, and auxiliary verbs, among others, which do not have a big impact on our predictions.
- c. From the Keras library, we used the Tokenizer function to create a dictionary of words based on their frequency. Then, we reduced the dictionary to the 1000 most repeated words and for each text we generated vectors with integer values that are the indexes in the dictionary. Finally, they were padded to the same length and entered the neural network described previously.

4.2 CNN

Convolutional neural networks (CNN) are a class of deep neural networks, most often applied to visual image analysis [26]. More recently, however, CNN have also found their place in solving problems related to NLP tasks. A CNN is generally build around the following layers:

- Convolutional layers are composed of neurons whose purpose is to detect patterns (features map) from their inputs.
- Pooling layers whose purpose is to reduce the feature map dimensionality in order to be more computationally efficient. These layers are often chained after convolutional layers.
- These two previous set of layers are generally followed by one or more fully connected layers

As previously stated, CNN has been a game changer in the field of image analysis. They also have shown some interesting results in NLP [18, 27]. Indeed, texts can be represented an array of vectors, just like images can be represented by an array of pixel values. Here, we deal with one dimensional convolutions, but the principles remain the same: we still want find patterns in the sequence which become more complex with each added convolutional layer.

In order to associate each word with a specific vector, different techniques can be used. We are talking here about words embedding techniques. The most effective ones are those that are context-sensitive. We can cite here Glove [28] and Word2Vec [29] and as we will see in the next section, Bert is another word embedding technique. Word embeddings are able, by reducing the dimension, to capture the context the semantic and syntactic similarity (gender, synonyms, , etc.) of a word. For example, one would expect the words “remarkable” and “admirable” to be represented by relatively closely spaced vectors in the vector space where these vectors are defined. It has been shown that contextual words embedding can effectively

capture the semantic and arithmetic properties of a word. It also reduces the size of the problem and therefore the learning task. In the case of CNN, the obtained inputs vectors will allow the model to have a much better representation of the words during the learning phase.

Given the sequential nature of texts, RNN and LSTM are more common models in NLP. However, they can be long and difficult to train. In this way, for large datasets CNN can be an interesting alternative.

Similar to the architecture selection procedure performed in LSTM, Table 4 shows the specifications of the three CNN architectures selected to subsequently perform an exhaustive search, varying the learning rate and batch size with the HyperOpt library. In addition, the text pre-processing performed before entering the CNN is also the applied in the LSTM Sect. 4.1.

4.3 Transformers

The sequential nature of RNNs was regularly pointed out as a hindrance to the training of these models on long texts both for computation time reasons (even modern GPUs do not parallelize well this type of process) and because of gradient vanishing problems (despite the use of LSTM models). In order to no longer process texts sequentially, transformers models provide a solution by processing the whole sequence all at once and by using the attention mechanism which allow capture different types of relationships between tokens. A transformer is built based on an encoder and a decoder, each of them consisting of a stack of attention and dense layers. Since 2017 and the first transformer model, many models have been developed such as ELMo, GPT-2 and GPT-3, BERT, XLNet, RoBERTa, and Turing-NLG. The main advantages of the transformer model are the following:

- The distance between two tokens is no longer a parameter taken into account by the model (the model can take into account long-term dependencies).
- The attention matrix calculation allows to parallelize the process of encoding and then decoding the sequences, thus accelerating the calculations.
- No labeled data are required to pre-train these models and it is then possible to train a transformer-based model by providing a huge amount of unlabeled text data.

From the last point it follows that it is possible to do transfer learning with this trained model in order to perform other NLP tasks like text classification, namely entity recognition, text generation, etc. This is how in June 2017, Google presents BERT (Bidirectional Encoder Representations from Transformers) [30]. It is a transformer composed of a suite of encoders only ($N = 12$ or 24 depending

Table 4 Defined architectures for CNN

Archi. 1	Archi. 2	Archi. 3
Input (200)	Input (200)	Input (200)
Embedding (200)	Embedding (200)	Embedding (200)
Conv1D (128, 3, ReLU)	Conv1D (16, 3, ReLU)	Conv1D (256, 3, ReLU)
MaxPooling1D (2)	Dropout (0.2)	Dropout (0.2)
Flatten ()	MaxPooling1D (2)	MaxPooling1D (4)
Dense (4, Sigmoid)	Conv1D (32, 3, ReLU)	Conv1D (300, 4, ReLU)
	Dropout (0.5)	Dropout (0.2)
	MaxPooling1D (2)	MaxPooling1D (4)
	Conv1D (64, 3, ReLU)	Conv1D (360, 4, ReLU)
	MaxPooling1D (2)	Dropout (0.5)
	Flatten ()	MaxPooling1D (4)
	Dense (4, Sigmoid)	Flatten
		Dense (400, ReLU)
		Dropout (0.2)
		Dense (4, Sigmoid)

on the version: base with 110 millions parameters or large with 340 millions parameters). Bert was originally pre-trained by using two tasks. It hides some of the words (15%, although this is actually more complex) and learns how to find them. This allows him to acquire a general and bidirectional knowledge of the language. BERT also learns to recognize if two sentences are consecutive or not. The corpus used for this pre-training was the BooksCorpus with 800 M words and a version of the English Wikipedia with 2,500 M words. When it came out, BERT was able to outperform state-of-the-art models for a large set of NLP tasks such as GLUE (General Language Understanding Evaluation) or SQuAD (Stanford Question Answer Dataset).

In the following, we use transfer learning (with our vigilance data) on two pre-trained models based on the BERT architecture and French corpus.

4.3.1 CamemBERT

CamemBERT [31] is based on RoBERTa, which improves the original implementation of BERT by using dynamic masking and training with larger batches and for longer. However, the main differences between RoBERTa and CamemBERT are that the latter uses the whole-word masking and SentencePiece tokenization instead of WordPiece. CamemBERT has been pre-trained on the French subcorpus (138GB) of a huge multilingual corpus (Common Crawl Oscar corpus) and proved that crawled data with high variability are preferable to using Wikipedia-based data. In the present article, we run the base implementation of CamemBERT from the HuggingFace's transformers library in Pytorch, which uses the original

architecture of BERT base (12 layers, 768 hidden dimensions, 12 attention heads, and 110M parameters).


4.3.2 FlauBERT

FlauBERT [32] is also a RoBERTa-based model for French. It has been pre-trained on less but more edited data (71GB). Its text corpus consists of 24 sub-corpora collected from different sources with diverse subjects and writing styles. Its performance is very close to CamemBERT. In fact, in its original paper, it is demonstrated that the performance of a parsing model can be improved with an ensemble of FlauBERT and CamemBERT. In the present article, we run the base implementation of FlauBERT from the HuggingFace's transformers library in Pytorch, which uses the original architecture of BERT base too.

5 Results and discussion

First, we present the scores of the univariate time series models for the prediction of the number of interventions. The models perform constant predictions equal to the mean, predictions corresponding to the persistence model (we predict the last known values), and predictions corresponding to the mean per hour (see Fig. 4). Next, we develop multivariate time series models, based on one of the most powerful learning machine tools available today: XGBoost [33]. The data used are composed of calendar features (day in the week, month, year, etc.), quantitative meteorological variables, corresponding to the three meteorological stations closest to the Doubs, which makes it possible to measure the extent to which predictions of firefighting interventions can be made by taking only

Fig. 4 Total person rescue interventions and its mean per hour series, early 2012



numerical data from meteorological information, and only vigilance levels from bulletins as indicators. Then, we reframe the problem as a multilabel classification task in order to analyze the impact of the tabular data previously described but with a new perspective that allows us to recognize the periods with intervention peaks and not the number of interventions. For this purpose, we compared the XGBoost classifier and the one from a technique also recognized for its speed and robustness, called random forest [34]. Finally, we replaced the tabular data with the meteorological texts of the bulletins and applied NLP techniques to study how valuable these data can be.

5.1 Prediction of interventions using basic univariate time series models

Scores corresponding to constant predictions are given in this subsection. The results obtained are shown in Table 5 for the Emergency Person Rescue, Table 6 for the Total Person Rescue, Table 7 for the Heating related interventions, and Table 8 for the Storm/Flood ones. In these tables, MAE stands for mean absolute error, and RMSE is for the root mean squared error, the most usual metrics for regression. As can be seen, the average per hour does better than the average alone for personal assistance, which is understandable given the daily seasonality of human activity. As shown in Fig. 4, there are fewer interventions at night than during the day because people are simply sleeping; similarly, because people eat at noon, there is a plateau at that time. This improvement is not found for

Table 5 Emergency person rescue prediction scores

Method	MAE	RMSE
Mean	1.9856	2.4554
Persistence	1.3504	1.8292
Mean/hour	1.6485	2.1160

Table 6 Total person rescue prediction scores

Method	MAE	RMSE
Mean	3.6939	4.5267
Persistence	2.0571	2.7298
Mean/hour	2.5843	3.3901

Table 7 Heating prediction scores

Method	MAE	RMSE
Mean	0.4829	0.6240
Persistence	0.1737	0.4490
Mean/hour	0.4641	0.6127

Table 8 Storm and flood prediction scores

Method	MAE	RMSE
Mean	0.4340	1.1742
Persistence	0.1749	0.5901
Mean/hour	0.4225	1.1699

interventions such as heating or storm, as these are only minimally related to human activity: A storm can cause damage both day and night.

One might a priori be astonished that the simplistic persistence model does better than the average per hour, which seems more evolved. However, this is well explained by the scarcity of interventions of the heating or storm type: If most of the time, there are 0 interventions per hour, then the probability that the slots at time t and time $t+1$ both have no intervention is high. Thus, replicating what happened at time t as a prediction of what will happen at time $t+1$ is a winning strategy when events are rare. The good success of this model in the case of rescue-type interventions is again explained by the strong daily seasonality of these interventions: Between one hour and the next, the number of interventions is close, when it is very

different between time t and $t+12$ hours. In other words, there is a strong correlation between the time series of the rescue-type and the time series shifted by one hour, as can be seen on the auto-correlation graph of the Total People Rescue, see Fig. 5.

5.2 Prediction of interventions using multivariate time series models

The first features to be taken into consideration are obviously the calendar data, which fully condition human activity. The time of day captures the daily seasonality, with a drop-in activity at night, a maximum of activity in the late morning and in the afternoon, with a plateau at noon. For various reasons, the day in the week also has its importance: weekend, a day without school for children, etc. In the same way, the day in the year makes it possible to recover particular periods such as the summer, the winter vacations, even particular days (national holiday, Christmas and New Year's Day, etc.). These particularities are also contained, but in a less continuous way, in the month in the year. Finally, the year is important because it allows us to model the general upward trend, for the various reasons mentioned above (aging population, disengagement of the private sector, etc.). For accidents related to heating (chimney fire, etc.), storms and floods, or even emergency rescue, it is reasonable to think that meteorology is important and that predictions should be improved by adding variables describing it. For example, between two national holidays, if the first one is rainy when the second one takes place under a radiant sun during a heat-wave, this will probably have an impact on firefighters' exits.

The purpose of this section is to see if this meteorological influence can be recovered without using NLP. We will therefore compare the quality of the predictions for the four types of intervention, with or without the temporal data (calendar), with or without the color of the weather

alert (vigilance), and with or without the quantitative data from the Météo-France site (weather) such as temperature, pressure, and wind. To do so, we will randomly separate our dataset between learning (80%) and test (20%), and we will look at the MAE and RMSE scores on the test set, once the learning is completed. This learning will be done with XGBoost (Poisson regression, max depth of 6), with 20% of the learning set used for validation, and an early stopping criterion of ten steps.

The first lesson to be learned from Table 9 is that, in the case of emergency person rescue, calendar data are obviously what appears to be most important, but that predictions can be improved by adding the color of the weather alert bulletin. These bulletins are basically too coarse a dataset to be useful in predicting these types of interventions on their own. As such, quantitative meteorological data do a little better but are far from what is obtained with calendar data. However, the best result is obtained by mixing calendar data with vigilance bulletins, when the addition of quantitative weather variables systematically lowers the quality of the predictions. The information that they carry on their own is found in a couple of calendar and vigilance data, which come as if cleaned of the noise that the quantitative variables carry: Variables such as gust over a period or mean wind speed (10 min.) have too local a scope, both spatially and temporally, whereas the bulletins have a more general scope (the information is digested). Note again that, among baselines, one does not do better there than the persistence model. An autoregressive component is fundamental to this prediction problem, and it would improve the scores of Table 9 in an obvious way: Among the interventions between t and $t+1$, there are the new interventions that appear after t , and those that appeared previously but are still in use (the system has strong inertia).

The same lessons can be learned from total interventions in a more pronounced way, see Table 10. This time, the best result is achieved by calendar data alone, and the

Fig. 5 Auto-correlation graph for total person rescue




Table 9 Prediction scores using XGBoost, emergency person rescue case

Features	MAE	RMSE
Calendar	1.523	1.961
Weather	1.825	2.307
Vigilance	2.021	2.493
Weather + calendar	1.586	2.049
Calendar + vigilance	1.400	1.881
Weather + vigilance	2.186	2.809
All	1.940	2.632

addition of weather variables only pushes XGBoost to overfit. These total interventions include, in addition to the emergency rescue, accidents on the public highway, and non-emergency rescue: person trapped in an elevator or locked in a balcony, wasp's nest, etc. These interventions are, by nature, much more difficult to predict. And the impact of temperature or atmospheric pressure is certainly much less on this type of intervention than, for example, knowing whether it is the middle of the day or the middle of the night. One can then naturally wonder if adding textual information on the weather could improve such results.

In the case of interventions related to heating, this time the best score is obtained by linking the calendar data to the vigilance bulletins, as shown in Table 11. This is understandable, given the nature of the interventions (chimney fire, electrical heating appliance fire, etc.). Here again, quantitative meteorological information does not provide the same benefit as the color of the vigilance bulletin, for the same reasons as previously mentioned. Conversely, for events such as storms and floods, if the calendar alone produces the best results, it is closely followed by the combination of calendar and meteorological data (cf. Table 12). Here, in a counterintuitive way, the vigilance bulletins greatly reduce performance. All this can be explained by noting first of all that floods do not occur at

Table 10 Prediction scores using XGBoost, total person rescue case

Features	MAE	RMSE
Calendar	2.223	2.917
Weather	3.281	4.138
Vigilance	4.117	5.138
Weather + calendar	2.375	3.118
Calendar + vigilance	2.279	2.998
Weather + vigilance	4.293	5.376
All	3.385	4.445

Table 11 Prediction scores using XGBoost, heating case

Features	MAE	RMSE
Calendar	0.413	0.558
Weather	0.502	0.663
Vigilance	0.510	0.659
Weather + calendar	0.475	0.636
Calendar + vigilance	0.319	0.495
Weather + vigilance	0.565	0.825
All	0.533	0.804

any time of the year in the Doubs, but mainly in winter. They are very localized in time, when the bulletins of vigilance usually extend over a fairly long period. Finally, such events follow a strong fall of water, a fact that is found in weather variables.

To conclude, the case of heating-type interventions shows that by adding information about the weather, predictions can be improved. The case of storms and floods, on the other hand, shows that quantitative information (temperature, pressure, etc.) and qualitative information (risk of storm, flooding, etc.) do not provide enough information on the weather: the latter has an obvious impact on the interventions, but these variables do not improve the predictions. And the Emergency and Total People Rescue confirm that these variables describe the weather forecast too crudely, which explains why we are interested in the textual content of the vigilance bulletins.

5.3 Prediction of intervention peaks using multilabel classification models based on decision trees and tabular data

In this section, the problem is restated as a multilabel classification, where the labels are the four categories Emergency Person Rescue, Total Person Rescue, Heating, and Storm/Flood, and the possible classes are 0 and 1 calculated as described in 3.2. This is in order to identify

Table 12 Prediction scores using XGBoost, storm/flood case

Features	MAE	RMSE
Calendar	0.325	0.734
Weather	0.386	0.867
Vigilance	1.729	4.959
Weather + calendar	0.370	0.863
Calendar + vigilance	0.793	2.453
Weather + vigilance	1.212	2.565
All	1.161	2.978

the peak periods of intervention concerning each category. The quantitative and qualitative variables described in the preceding section are the inputs for the models developed in this section, where each sample of the dataset represents one hour, an illustration is shown in Fig. 6. In addition, the models created are the results of various combinations of the variables to observe the influence of them but with a classification approach. These models will be the baselines to be overcome by the NLP techniques in the following section.

Before applying NLP techniques, we also sought to develop and compare models with simpler techniques, which do not require a greater consumption of resources, such as those based on decision trees, but which have demonstrated their remarkable effectiveness in the literature with tabular data. We chose XGBoost as a representative of boosting algorithms, which seek to reduce model variance, and random forest, as a representative of bagging algorithms, which seek to reduce bias.

Since decision trees are robust to work with categorical and continuous values, we kept the original values of our variables. Furthermore, similar to the previous section, we randomly split the data into 80% for training and 20% for testing. Since XGBoost and Random Forest do not natively support multitarget classification, we used *MultiOutputClassifier*, from the Scikit-Learn library, to fit one classifier per target. To select the best model, we used Bayesian optimization via the HyperOpt library. In total 100 iterations were performed for each combination of variable type and each technique. To guide the search for the best model, we set as loss function the metric *Micro F1-score*, hereafter called *F1-score* for short. Generally, this metric is used to assess the quality of multilabel binary problems, where the score closer to 1 means better *Micro-Precision* and *Micro-Recall* (we will abbreviate both to *Precision* and *Recall*, respectively) and closer to 0 means poor model performance. Other metrics such as *Accuracy* and *Balanced Accuracy* are also presented to analyze our resulting models.

Hour	Day	Weekday	...	Emergency Person Rescue	Total Person Rescue	Heating	Storm/Flood
...
0	4	2	...	1	0	1	0
1	4	2	...	1	0	0	1
...
10	22	6	...	0	1	1	0
11	22	6	...	1	1	0	0
...

Fig. 6 Illustration example of the tabular data for multilabel classification, considering calendar variables

Table 13 shows the results of the best models obtained for each data input combination and Table 16 describes their hyperparameters and the search space used. As we can see in Table 13, calendar data together with vigilance alert levels improve the performance of the models. What is more, the best model, obtained by random forest, used only calendar data reaching an F1-score of 0.81. Also, we note that when the inputs are weather, vigilance indicators, and weather plus vigilance indicators, the models show an F1-score below 0.64, a poor performance, when in fact these variables should present a greater contribution on the recognition of intervention peaks.

Therefore, these are the baseline results that we need to outperform to be efficient. An NLP tool must have an F1-score greater than those values. Thus, we have to see whether information derived from the text of the bulletins allows better prediction of interventions than simple calendar data, vigilance levels, or quantitative information such as temperature or wind speed. For this purpose, we are interested in comparing different NLP models.

5.4 Prediction of intervention peaks using multilabel classification models based on NLP techniques and text from meteorological bulletins

In this section, we seek to discover if we have better learning for the four outputs when we process the texts of the bulletins.

For this task, the dataset used considers the bulletins of the three neighboring departments of Doubs and the bulletins of Doubs. They were structured as shown in Table 2 and the task is maintained as a multilabel classification for the recognition of the intervention peak periods.

The initial dataset was split into 80% for the learning phase and 20% for the testing phase. Experimentation was then performed by varying different hyperparameters for the different models. The results presented in Table 14 correspond to the best results obtained for each technique applied. We calculated the same metrics presented in previous Sect. 5.3 to identify possible improvements. On the one hand, as described in Sects. 4.1 and 4.2, for LSTM and CNN, we applied a text preprocessing and used the library Hyperopt to search for the best configuration for our models, where the number of iterations was 100, and the guiding metric was the F1-score. The best setting for LSTM was the architecture no. 1, with a learning rate = $6e-4$, number of epochs = 200, and the batch size = 59. The best configuration for CNN was the architecture no. 3 with a learning rate = $9e-3$, number of epochs = 105, and the batch size = 95. On the other hand, transformers were used in its base models: CamemBERT (110 Million parameters) and FlauBERT (138 Million parameters) with

Table 13 Prediction results of the multilabel models based on XGBoost and Random Forest techniques

Technique	Input	F1-score	Accuracy	Balanced accuracy	Precision	Recall
XGBoost	Calendar	0.80	0.48	0.82	0.80	0.80
	Weather	0.63	0.21	0.64	0.65	0.60
	Vigilance	0.62	0.23	0.62	0.61	0.62
	Weather + calendar	0.71	0.30	0.71	0.73	0.69
	Calendar + vigilance	0.81	0.48	0.82	0.81	0.80
	Weather + vigilance	0.64	0.24	0.65	0.65	0.64
	All	0.71	0.35	0.71	0.74	0.68
Random Forest	Calendar	0.81	0.51	0.83	0.82	0.81
	Weather	0.61	0.24	0.61	0.63	0.59
	Vigilance	0.64	0.18	0.58	0.54	0.77
	Weather + calendar	0.73	0.35	0.72	0.74	0.71
	Calendar + vigilance	0.81	0.50	0.83	0.81	0.81
	Weather + vigilance	0.62	0.24	0.61	0.63	0.61
	All	0.72	0.33	0.72	0.74	0.71

Table 14 Prediction results of the multilabel models based on NLP techniques

Technique	Input	F1-score	Accuracy	Balanced accuracy	Precision	Recall
CNN	Bulletin text	0.84	0.56	0.78	0.83	0.85
LSTM		0.85	0.56	0.80	0.84	0.86
FlauBERT		0.87	0.59	0.82	0.86	0.89
CamemBERT		0.89	0.65	0.84	0.87	0.90

the following hyperparameters. For CamemBERT: learning rate = $1e-5$, number of epochs = 75, and batch size = 48. For FlauBERT: learning rate = $1e-5$, number of epochs = 150, and batch size = 128. The search of the best hyperparameters was performed by using a random search with an early stopping strategy of 15 iterations. For more details on the description of the search space and the best configuration, see Table 17.

From Table 14, we see that LSTM and CNN results are quite similar. However, the two French transformers models, CamemBERT and FlauBERT, outperform both traditional techniques. This is mainly because the LSTM and CNN models were trained only with the vocabulary of the bulletins while the transformers were pre-trained with an extensive vocabulary of the French language (Common Crawl Oscar corpus). Thus, this set of experiments confirm the recent literature results [35] on text classification problems and the superiority of transformers models. Note that, Accuracy and Balanced Accuracy are quite different due to the imbalance of the dataset as mentioned in Sect. 3.2. Nevertheless, when looking at the F1-score, all the models in this section remarkably outperform the results obtained in the previous section with decision trees (Table 13). What is more, the best model obtained with CamemBERT overcomes by far the best model obtained with random forest by 8%, 14%, 1%, 5%, and 9%, when

comparing F1-score, Accuracy, Balanced Accuracy, Precision, and Recall, respectively.

When we examine accuracy by independent category in Table 15, we prove that the feature extraction from texts improves the recognition of intervention peaks due to rare events. The best NLP model reached accuracies of 80%, 86%, 92%, and 86% for Emergency Rescue People, Total Rescue People, Heating, and Storm/Flood, respectively. Moreover, the last two categories Heating and Storm/Flood, which represent interventions generated by rare events and which were complicated to recognize by the approaches analyzed in the previous sections, are the ones that demonstrated high accuracy with the NLP techniques and bulletin texts, without degrading the accuracy of the two rescue-type categories.

As mentioned in Sect. 2, generally, the models developed for predicting the interventions of fire departments or EMS use tabular data such as temporal information, quantitative meteorological variables, and traffic indicators [3, 5, 6, 11]. These are very useful for recognizing incidents related to human activity since it is possible to identify seasonality and trend over time. For example, people are more active during the day than at night, there are more drownings in the pool during the summer than in the winter, as the population increases the incidents also increase, etc. Also, because the major operational burden of these organizations is the interventions rescue-type.

Table 15 Accuracy results for each type of intervention, considering the models generated with the weather bulletins

Model	Emergency person rescue	Total person rescue	Heating	Storm/flood
CNN	0.76	0.82	0.84	0.82
LSTM	0.76	0.83	0.85	0.83
FlauBERT	0.79	0.87	0.88	0.85
CamemBERT	0.80	0.86	0.92	0.86

However, some interventions are difficult to detect, since they occur only a few times a year. These are produced by rare events such as natural phenomena (storms, floods, forest fires, etc.), and although their occurrence is minimal over the years, the workload they produce in a small period can be 28 times more than normal in some cases. For example, in 2016, the average number of interventions assisted by SDIS 25 per hour was 3.34. However, there was an hour in which 84 interventions occurred due to a storm. The storm caused flooding in the region, human and material damage, and breakdowns in the service of SDIS 25 [3]. For this reason, we need an intelligent system that could help to predict the peak periods of intervention generated by rare events. Thus, the present work developed models based on NLP techniques and meteorological bulletins from public sources to recognize periods with a heavy operational load. The results obtained are significant for practical purposes. Initially, the model could be deployed in production as a small stand-alone application. Or, the predictions (the binary indicators) could be included in a larger set of tabular data that would be the input for a predictive model of the number of interventions for a certain time and location. In this way, with an initial application or with a more robust system, the fire department could reorganize its personnel and armament to cope with these periods of high demand, reduce breakdowns in service due to lack of resources, and save more lives.

6 Conclusion

The present paper demonstrates the effectiveness of NLP techniques for the recognition of rare events that will cause an increase above the average in certain periods of firefighter interventions. This is done by processing the texts

contained in the weather bulletins using the traditional techniques LSTM and CNN, and transformers CamemBERT and FlauBERT. The results of the NLP models and bulletin texts exceed those of the baselines with Decision Trees (XgBoost and Random Forest) and tabular data by 8% and 14% when comparing the best F1-score and Accuracy, respectively. The advantage of using these texts is also reflected when assessing the accuracy of the two categories with interventions related to rare events, achieving 92% for Heating and 86% for Storm/Flood with the best CamemBERT model developed.

In this way, fire departments and EMS, in general, would be able to identify peak periods of interventions and optimize their response by establishing better strategies to prepare their armament for natural disasters (storms, floods, etc.) and keep the population better protected and safe.

As future work, we propose to add meteorological bulletins from other departments of the country, which would allow us to better track a possible extreme event and its consequences on the firefighters' workload. Furthermore, we will develop and compare other modeling and text preprocessing techniques with the texts in French and English. Finally, we aim to integrate the results of this classification approach into a larger regression model that predicts the number of interventions for a certain period (hourly, daily, and monthly) and by locality (principal cities and mountain cities).

Appendix

Settings used during modeling

See Tables 16 and 17.

Table 16 Hyperparameters search space and the best configuration for XGBoost and Random Forest multilabel models

Method	Search Space	Best configuration						
		Calendar	Weather	Vigilance	Weather + Calendar	Calendar + Vigilance	Weather + Vigilance	All
XGBoost	n_estimators: [50–200]	195	159	193	132	132	86	112
	learning_rate: [0.001–0.8]	0.49	0.14	0.39	0.17	0.40	0.60	0.69
	max_depth: [1–100]	100	6	2	3	20	1	4
	colsample_bytree: [0.2–1]	0.99	0.5	0.56	0.42	0.87	0.22	0.46
	objective: multi:softmax eval_metric: mlogloss	multi:softmax mlogloss						
Random Forest	n_estimators: [50–500]	52	328	61	152	411	337	87
	max_features: [0.2–1]	0.38	0.74	0.88	0.46	0.53	0.94	0.22
	max_depth: [1–10]	100	1	10	5	20	1	5
	class_weight: [balanced, balanced_subsample]	Balanced_subsample	Balanced	Balanced	Balanced	Balanced	Balanced	Balanced

Table 17 Hyperparameters search space and the best configuration for NLP multilabel models

Method	Search space	Best configuration
CNN	Type of architecture: [1,2,3]	3
	Learning rate: [0.00001–0.01]	0.009
	Batch size: [40–150]	95
	Epochs: 500	105
	Early stopping: 15	15
	Restore best weights: True	True
LSTM	Type of architecture: [1,2,3]	1
	Learning rate: [0.00001–0.009]	0.0006
	Batch size: [40–150]	59
	Epochs: 200	200
	Early stopping: 20	20
	Restore best weights: True	True
FlauBERT	Type of architecture: flaubert-base-cased	Flaubert-base-cased
	Learning rate: [0.0001, 0.00001]	0.00001
	Batch size: [16–256]	128
	Epochs: [10–200]	150
	Early stopping: 15	15
	Restore best weights: True	True
CamemBERT	Type of architecture: camembert-base	Camembert-base
	Learning rate: [0.0001, 0.00001]	0.00001
	Batch size: [16–256]	48
	Epochs: [10–200]	75
	Early stopping: 15	15
	Restore best weights: True	True

Acknowledgements This work was funded by the Service Départemental d’Incendie et de Secours du Doubs (SDIS 25), with the support of the French Ministry of Higher Education and Research (managed by the National Association of Research and Technology (ANRT) for the CIFRE thesis (N 2019/0372), and by the EIPHI-BFC Graduate School (contract “ANR-17-EURE-0002”). Computations have been performed on the supercomputer facilities of “Mésocentre de Calcul de Franche-Comté.”

Declarations

Conflict of interest Authors certify that there is no conflicts of interests or competing in interests in these works.

References

- Cerna S, Guyeux C, Royer G, Chevallier C, Plumerel G (2020) Predicting fire brigades operational breakdowns: a real case study. *Mathematics* 8(8):1383
- Arcolezi HH, Couchot J-F, Cerna S, Guyeux C, Royer G, Bouna BA, Xiao X (2020) Forecasting the number of firefighter interventions per region with local-differential-privacy-based data. *Comput Secur* 96:101888
- Cerna S, Guyeux C, Arcolezi HH, Couturier R, Royer G (2020) A comparison of LSTM and XGBoost for predicting firemen interventions. In: *Trends and innovations in information systems and technologies*. WorldCIST 2020. Springer International Publishing, vol. 1160, pp. 424–434
- Pirklbauer K, Findling RD (2019) Predicting the category of fire department operations. In: *Proceedings of the 21st international conference on information integration and web-based applications & services*. ACM
- Morello TF, Ramos RM, Anderson LO, Owen N, Rosan TM, Steil L (2020) Predicting fires for policy making: improving accuracy of fire brigade allocation in the Brazilian amazon. *Ecol Econ* 169:106501
- Sankar SH, Jayadev K, Suraj B, Aparna P (2016) A comprehensive solution to road traffic accident detection and ambulance management. In: *2016 international conference on advances in electrical, electronic and systems engineering (ICAEES)*. IEEE, Nov. [Online]. Available: <https://doi.org/10.1109/icaees.2016.7888006>
- Chen AY, Lu T-Y, Ma MH-M, Sun W-Z (2016) Demand forecast using data analytics for the preallocation of ambulances. *IEEE J Biomed Health Inform* 20(4):1178–1187
- de la Mota IF, Perez ES, Garduno AV (2017) Optimization and simulation of an ambulance location problem. In: *2017 winter simulation conference (WSC)*. IEEE, Dec. [Online]. Available: <https://doi.org/10.1109/wsc.2017.8248188>
- Carvalho A, Captivo M, Marques I (2020) Integrating the ambulance dispatching and relocation problems to maximize system's preparedness. *Eur J Oper Res* 283(3):1064–1080
- Rasel RI, Sultana N, Islam GA, Islam M, Meesad P (2019) Spatio-temporal seismic data analysis for predicting earthquake: Bangladesh perspective. In: *2019 research, invention, and innovation congress (RI2C)*, IEEE
- Lu XS, Zhou M, Qi L (2017) Analyzing temporal-spatial evolution of rare events by using social media data. In: *2017 IEEE international conference on systems, man, and cybernetics (SMC)*, IEEE
- Cook BL, Progovac AM, Chen P, Mullin B, Hou S, Baca-Garcia E (2016) Novel use of natural language processing (NLP) to predict suicidal ideation and psychiatric symptoms in a text-based mental health intervention in madrid. *Comput Math Methods Med* 2016:1–8
- Hasan MR, Maliha M, Arifuzzaman M (2019) Sentiment analysis with NLP on twitter data. In: *2019 international conference on computer, communication, chemical, materials and electronic engineering (IC4ME2)*, IEEE 2019
- Ding Y, Li B, Zhao Y, Cheng C (2017) Scoring tourist attractions based on sentiment lexicon. In: *2017 IEEE 2nd advanced information technology, electronic and automation control conference (IAEAC)*, IEEE 2017
- Kedia A, Rasu M (2020) Hands-on Python natural language processing: explore tools and techniques to analyze and process text with a view to building real-world NLP applications. Packt Publishing Ltd
- Preum S, Shu S, Hotaki M, Williams R, Stankovic J, Alemzadeh H (2019) CognitiveEMS. *ACM SIGBED Rev* 16(2):51–60
- Temnikova I, Castillo C, Vieweg S (2015) Emterms 1.0: a terminological resource for crisis tweets. In: *ISCRAM*
- Kozlowski D, Lannelongue E, Saudemont F, Benamara F, Mari A, Moriceau V, Boumadane A (2020) A three-level classification of french tweets in ecological crises. *Inf Process Manag* 57(5):102284
- Karami A, Shah V, Vaezi R, Bansal A (2019) Twitter speaks: a case of national disaster situational awareness. *J Inf Sci* 46(3):313–324
- Corvey WJ, Vieweg S, Rood T, Palmer M (2010) Twitter in mass emergency: what nlp techniques can contribute. In: *HLT-NAACL 2010*
- Fersini E, Messina E, Pozzi FA (2016) Earthquake management: a decision support system based on natural language processing. *J Ambient Intell Humaniz Comput* 8(1):37–45
- Sermet Y, Demir I (2018) An intelligent system on knowledge generation and communication about flooding. *Environ Model Softw* 108:51–60
- “Météo-France public data,” https://donneespubliques.meteo-france.fr/?fond=produit&id_produit=90&id_rubrique=32, accessed: 2021-02-01
- “Vigilance cards and bulletins archive, Météo-France,” <http://vigilance-public.meteo.fr/index.php>, accessed: 2021-02-01
- Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Comput* 9(8):1735–1780
- Goodfellow I, Bengio Y, Courville A (2016) *Deep learning*. MIT press, Cambridge
- Nguyen DT, Mannai KAA, Joty S, Sajjad H, Imran M, Mitra P (2016) Rapid classification of crisis-related data on social networks using convolutional neural networks. In *CoRR*, vol abs/1608.03902. <http://arxiv.org/abs/1608.03902>
- Pennington J, Socher R, Manning CD (2014) Glove: global vectors for word representation. In: *Empirical methods in natural language processing (EMNLP)*, pp. 1532–1543. [Online]. Available: <http://www.aclweb.org/anthology/D14-1162>
- Mikolov T, Chen K, Corrado G, Dean J (2013) Efficient estimation of word representations in vector space. In: *1st Int. Conf. on Learning Representations*, Scottsdale, Arizona, USA. <http://arxiv.org/abs/1301.3781>
- Devlin J, Chang M-W, Lee K, Toutanova K (2019) Bert: pre-training of deep bidirectional transformers for language understanding. In *CoRR*. vol. abs/1810.04805. <http://arxiv.org/abs/1810.04805>
- Martin L, Müller B, Suárez PJO, Dupont Y, Romary L, de la Clergerie ÉV, Seddah D, Sagot B (2019) Camembert: a tasty french language model. In: *CoRR*, vol. abs/1911.03894. [Online]. Available: <http://arxiv.org/abs/1911.03894>
- Le H, Vial L, Frej J, Segonne V, Coavoux M, Lecouteux B, Allauzen A, Crabbé B, Besacier L, Schwab D (2019) Flaubert: unsupervised language model pre-training for french. In: *CoRR*, vol. abs/1912.05372. [Online]. Available: <http://arxiv.org/abs/1912.05372>
- Chen T, Guestrin C (2016) Xgboost. In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, Aug 2016
- Breiman L (2001) Random forests. *Mach Learn* 45(1):5–32
- Minaee S, Kalchbrenner N, Cambria E, Nikzad N, Chenaghlu M, Gao J (2021) Deep learning based text classification: a comprehensive review. In *ACM Computing Surveys* 54(3):1–40. <https://doi.org/10.1145/3439726>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.