

Traditional vs. AI-Generated Meteorological Risks for Emergency Predictions

Naoufal Sirri¹ and Christophe Guyeux¹

FEMTO-ST Institute, UMR 6174 CNRS, University of Franche-Comté, Belfort,
France, {naoufal.sirri,christophe.guyeux}@univ-fcomte.fr

Abstract. This study aims to analyze and examine in-depth the feature selection process using Large Language Models (LLMs) to optimize firefighter prediction performance. Although features from reliable sources are known to significantly aid predictions, their accuracy may be limited in critical situations requiring rigorous prioritization. Therefore, the focus was placed on meteorological risks for a comparative diagnosis between their extraction from Météo France and those generated by LLMs across various dimensions. Given the crucial role of meteorological risks as key informational sources for decision-making, this study explores the impact of feature extraction methods related to these risks on predicting firefighter interventions over nine years, from 2015 to 2024. Annual reports on firefighter activities in France highlight the growing influence of weather-related risks, underscoring the urgent need for precise and actionable meteorological information to support rapid and effective emergency response strategies. The methodology implemented involved comprehensive data preparation, an in-depth analysis of feature extraction through different approaches, and their evaluation from multiple perspectives. This required leveraging machine learning models such as XGBoost, Random Forest, and Support Vector Machines (SVM) to assess and analyze prediction results based on two feature spaces: F1 (including general features and meteorological risks extracted from Météo France) and F2 (including general features and meteorological risks generated by LLMs). The results revealed that models trained with the F2 feature space consistently demonstrated superior performance. Notably, annual improvements were observed, particularly for high and very high intervention activities. However, the use of the F2 space proved less effective for low intervention activities and underperformed compared to F1 during the summer season. In conclusion, this work presents a concrete methodology for forecasting and enhancing resource management, accelerating firefighter response times, and ultimately contributing to life preservation by reducing the risk of failure during critical incidents.

Keywords: Firefighters intervention · Feature selection · Prediction · XGBoost · Large language model (LLM).

1 Introduction

The use of artificial intelligence to anticipate the needs of emergency services has become a crucial tool for managing urgent situations and accurately forecasting operational requirements. This is particularly important in sectors like environmental protection and public safety. Enhancing the predictive performance of firefighter interventions and optimizing resource allocation relies on extracting features for machine learning models from reliable and precise data sources. This is particularly relevant for meteorological risk features, where inaccurate estimations can lead to serious consequences, such as insufficient resource deployment and delayed response times. Assessing features extracted from reliable sources over a nine-year period (2015–2024) to develop robust predictions presents significant challenges. Using diverse approaches to identify meteorological risk features, whether derived from traditional sources such as Météo France or generated by Large Language Models (LLMs) [1], this study aims to evaluate the outcomes from multiple perspectives. The objective is to select the most effective feature set for predictive modeling. The objective is to identify the best feature set for predictive modeling. These approaches facilitate a thorough comparative analysis and inform decision-making regarding the most suitable feature extraction method based on seasonal variations or firefighter activity levels.

Previous research has made significant strides in feature extraction and selection for prediction tasks across various domains, including fire intervention management. Recent studies have taken this further by employing generative AI to identify relevant features. For instance, [2] proposed a hybrid approach combining feature selection and extraction to enhance cancer prediction using DNA methylation data, effectively addressing challenges related to high dimensionality and noise. Similarly, [3] demonstrated that randomly reducing the number of instances before feature reduction could maintain high classification performance while minimizing data requirements and computational time. On the other hand, [4] highlighted the critical role of context selection in improving the accuracy of Large Language Models (LLMs) for data extraction, outperforming traditional methods and reducing errors. Additionally, [5] introduced an effective method for selecting representative samples for few-shot fine-tuning of LLMs, achieving enhanced performance with only 2% of the data and reducing computational costs by 97%.

However, limited research has focused on predicting firefighter interventions, and no study to date has specifically explored the application of generative AI, particularly Large Language Models (LLMs), in this field. For instance, [6] employed machine learning techniques to enhance resilience and efficiency through optimized firefighter deployment strategies in France. Similarly, [7] proposed leveraging natural language processing techniques to extract features from weather bulletins, achieving high accuracy (80-92%) in predicting intervention peaks caused by rare events. Finally, in relation to the topic addressed in this paper, [8] assessed the reliability of predictions concerning firefighter interventions, aiming to optimize resource utilization and response times. They also conducted several separate studies analyzing the impact of specific factors such

as air quality, solar activity, river height, and weather conditions on fire intervention activities, as detailed in [9], [10], [11], and [12], respectively.

Our research team is fully committed to this initiative by adopting an innovative approach: analyzing and evaluating meteorological risk features using various methodologies, including extraction from traditional sources such as Météo-France and data generated by Large Language Models (LLMs). This analysis is enriched by comprehensive historical datasets, offering a wealth of information on various aspects of firefighter interventions. The primary aim of this study is to assess the impact of feature extraction methods on improving the accuracy of firefighter intervention predictions. To achieve this, a hypothesis and its corresponding prediction are presented as part of this research effort:

- Hypothesis: Integrating meteorological risks generated by LLMs into predictive modeling provides a superior alternative to traditional sources (e.g., Météo-France) for optimizing firefighter intervention predictions.
- Prediction: Models incorporating features generated by LLMs will outperform those utilizing meteorological risks from traditional sources regarding predictive accuracy, especially for high-intervention activities.

To ensure the precision of this study, a well-structured methodological framework was established, meticulously adhering to its nuanced details as outlined below. Section 2 provides a comprehensive overview of the methodologies and materials utilized across the various experiments. The study’s progression is detailed in Section 3, which carefully presents the results obtained through this research. Section 4 delves into a thorough discussion, analyzing the findings in depth to address the initial research question. This section offers an extensive evaluation and diagnosis of the outcomes, emphasizing relevant implications and interpretations. Finally, Section 5 concludes the study by summarizing key discoveries, highlighting significant contributions, and proposing avenues for future research. This methodological framework facilitated a systematic exploration of generative AI’s capabilities in extracting meteorological risk features. Through a comprehensive comparison of two approaches from multiple perspectives, this analysis remained well-organized, rigorous, and critically focused throughout the investigation.

2 Material and Methods

2.1 Data Handling

Sources and Data Extraction In this study, the data collection phase played a crucial role. A comprehensive dataset was used, provided by the Fire and Rescue Services of Doubs, France. This dataset encompasses 334,536 intervention records spanning from January 1, 2015, to October 31, 2024. Each entry includes details such as start and end timestamps, location, duration, and type of intervention. To enrich this dataset with potentially relevant features, variables were collected from diverse sources.

- **General Features:** Basic features included school periods, holidays, solar activity, air quality data, and epidemiological information. Notably, epidemiological data were collected on diseases such as influenza, chickenpox, and acute diarrhea, sourced from the Sentinel Network [13]. Astronomical features, including lunar phases, moonrise, and distance between the Earth and the Moon, were added to evaluate their impact on night visibility [14, 15]. Satellite data from NASA’s VIIRS and MODIS missions were also integrated, leveraging different wavelengths and resolutions to assess wildfire propagation in specific regions [16, 17]. Calendar-related features captured details such as holidays, academic breaks [18], and time-related attributes like hours, days, weeks, months, and years. Additionally, solar activity metrics, such as sunspot area, the 10 cm radio flux, and the sunspot count, were included [19]. Air quality variables, including fine particles, ozone, and nitrogen oxides, were obtained from monitoring stations in the region [20].
- **Meteorological Features:** In a second phase, meteorological vigilance bulletins from Météo France covering the Doubs region were included. These reports detailed various weather risks with assigned color codes and included textual bulletins to enhance the overall meteorological understanding (refer to Table 1 and Fig. 1) [21].

The complete dataset was systematically enriched, integrating firefighter intervention data with these diverse features from multiple sources (refer to Table 2). This comprehensive dictionary serves as a robust foundation for the analysis conducted in this study.

Table 1: Categorization of meteorological risk by alert levels

Weather risk	Green	Yellow	Orange	Red
Wind	76544	1204	2602	0
Precipitations	75769	1560	3021	0
Thunderstorms	72863	3853	2745	889
Snow-Ice	79885	0	465	0
Heatwave	79061	324	965	0
Extreme Cold	80248	0	102	0

Data Refinement and Processing To optimize the input data for the machine learning model, several pre-processing techniques were applied. First, the "StandardScaler" technique from the Scikit-learn library [22] was employed to normalize numerical variables. This method adjusts the data distribution to achieve a mean of zero and a unit variance, ensuring consistent scaling across variables. The normalized features included the year, time of day, and statistics for diseases such as chickenpox, influenza, acute diarrhea, and the Earth-Moon distance. In addition, the TargetEncoder technique [23] was applied to encode categorical variables. This approach replaces each category with the mean of the

{**Texte_situation_actuelle**: 'Une profonde dépression se creuse en ce moment sur le Cotentin. Au sud de cette dépression, le vent continue de se renforcer sensiblement, avec des rafales comprises généralement entre 80 et 100 km/h sur le nord de la Bourgogne, la Champagne-Ardenne, la Lorraine et maintenant la Franche-Comté. Très localement des pointes entre 100 et 110 km/h ont été relevées.'

Texte_qualification: 'Nouvel épisode de vent violent nécessitant une vigilance particulière.'

Texte_situation_future: 'Actuellement, la dépression se dirige rapidement vers l'ouest de la Belgique. Dans la partie sud de sa trajectoire, un renforcement notable du vent est en cours et va persister cet après-midi et jusqu'en fin de journée sur les départements les plus à l'est. Au plus fort de l'épisode venteux (courant de l'après-midi), on attend généralement des rafales situées entre 90 et 100 km/h en plaine, voire localement et temporairement 100 à 110 km/h. Dans les zones à proximité immédiate du relief des Vosges, des rafales comprises entre 100 à 130 km/h sont également possibles. Il faudra attendre le début de nuit prochaine pour voir le vent nettement faiblir.'}

Fig. 1: Example of a weather alert bulletin (in French)

Table 2: Distribution of variables across feature families

Family of features	Number of variables
Calendar	9
Temporal	6
Holidays and academic breaks	8
Epidemiological	4
Astronomical	8
Solar activity	12
Air quality	11
Earth images	2
Meteorological	6

corresponding target variable. Features such as the day of the week, year, month, and public holidays were encoded using this method. Additionally, meteorological risk features were enriched by employing OpenAI's large language model (LLM). The LLM was solicited through a specific prompt designed to analyze weather vigilance bulletins and extract structured color-coded risk information. This prompt enabled the model to interpret textual descriptions and map them to predefined risk categories. The integration of these AI-generated features significantly enhanced the dataset by providing additional, highly relevant variables (refer to Fig. 2).

To facilitate effective classification, the original target variable representing the number of emergency interventions was transformed into an ordinal qualitative scale ranging from 0 to 3, based on percentile thresholds:

- 0: Low intervention activity
- 1: Moderate intervention activity
- 2: High intervention activity
- 3: Very high intervention activity

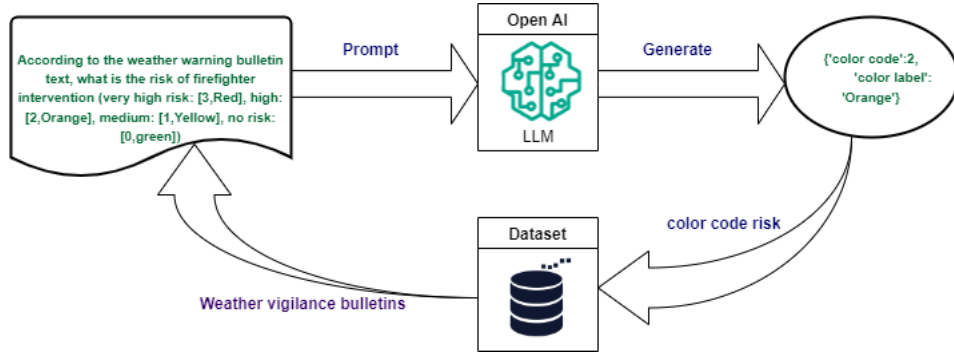


Fig. 2: Integration of AI-Driven features for enriching meteorological risk data

Exploratory Data Insights A thorough analysis of the dataset, extracted from multiple sources, is crucial to uncover relevant observations that can provide insights for this research. On average, 32,000 fire intervention events are recorded annually, equating to more than four interventions per hour. It was observed that the number of interventions has been consistently increasing each year due to various factors such as demographic aging, significant population growth, the effects of climate change, and the closure of small hospitals.

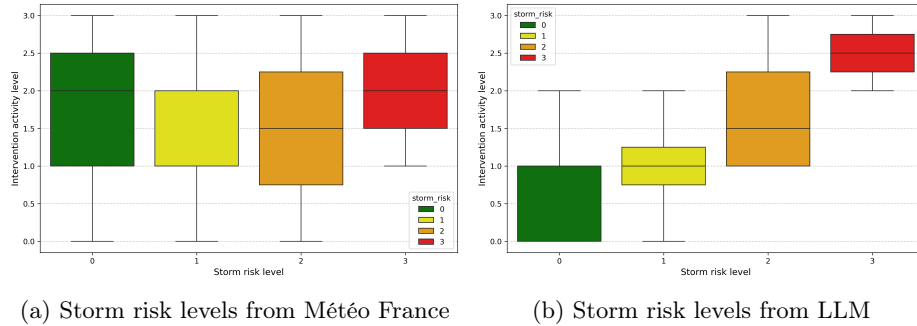


Fig. 3: Intervention activity distribution across storm risk levels

In this study, the analysis of meteorological risks played a vital role in identifying and optimizing key features for predicting firefighter interventions. For instance, Figures 3-(a) and 3-(b) illustrate the distribution of storm risks from Météo France compared to those generated by the Large Language Model (LLM). It is evident that the risk levels from Météo France are largely aligned with those of the LLM, which shows more dispersed values. This discrepancy is likely due to the fact that Météo France uses calibrated thresholds based on meteorological expertise, while the LLM-derived risks are based on text analysis, leading to

a wider variability due to linguistic nuances and the absence of direct calibration. Additionally, Figure 4 visualizes the average storm risks from Météo France which is below the LLM-derived risks in relation to firefighter activity. This difference likely arises because Météo France employs standardized meteorological models and thresholds for risk evaluation, whereas LLM risks are influenced by nuanced textual data, leading to broader and less conservative estimates.

In conclusion, feature evaluation from both Météo France and LLM regarding firefighter activity, assessed through feature importance analysis via XGBoost, revealed that features such as storms, wind, snow, and extreme cold from the LLM source were more significant than those from Météo France. Conversely, features like heatwaves and precipitation from Météo France were found to be more significant. This could be attributed to the fact that textual information related to heatwaves and precipitation may be less rich or more immediately interpretable compared to other meteorological risks.

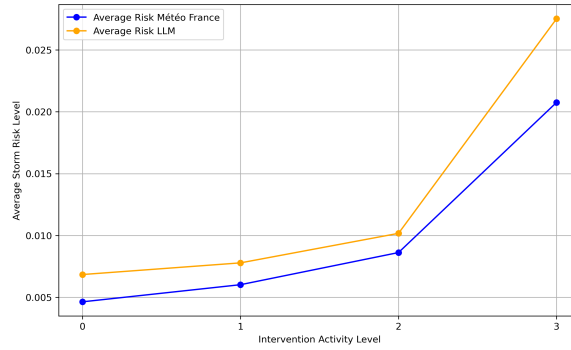


Fig. 4: Average storm risk levels: Météo France vs. LLM by intervention activity

2.2 Feature Engineering and Predictive Models

Feature Engineering Traditionally, all features in a dataset are included during model training, assuming each feature significantly contributes to building an optimal model. However, including all features has significant drawbacks. Some may exhibit strong correlations with others, while some may have limited usefulness. This can lead to suboptimal generalization or the inclusion of redundant information. Additionally, using all features can increase computational time without necessarily improving model performance [24]. Therefore, selecting a limited or optimized set of features can improve model performance. In this study, the "feature importance" technique was employed, which assigns a score to each variable based on its relevance to the target variable [25]. However, this technique was applied only to the general features. Meteorological risk features were analyzed separately in two feature spaces:

- F1: General features + Météo-France risk features.
- F2: General features + LLM-generated risk features.

This study used the following feature selection techniques:

1. Pearson and Kendall Correlation Coefficients: features with an absolute correlation of 0.4 or less with the target variable were excluded.
2. Chi-Square selector: the Chi-Square test was used to evaluate the dependency between each feature and the target variable, and then select the top "K" best variables.
3. Light Gradient Boosting model (LightGBM [26]): this model was employed with specified hyperparameters (learning rate = 0.1, objective = regression, metric = RMSE, number of leaves = 2^9 , maximum depth = 9, and number of estimators = 120,000). Early stopping was implemented after 10 iterations.

The final feature selection process produced a coherent set of features, emphasizing those identified as important across most methodologies and those with high scores.

Predictive Models To achieve meaningful results, three models were selected: Random Forest (RF) [27], Extreme Gradient Boosting (XGBoost) model [28], and Support Vector Machine (SVM) [29]. These models were chosen for their robustness, ability to handle complex datasets, and widespread use in classification problems. The combination of these models leverages their complementary strengths: XGBoost for capturing complex relationships, Random Forest for its resistance to noise, and SVM for its precision on well-separated data.

1. Random Forest (RF): Random Forest is an ensemble learning method that combines multiple decision trees, each trained on random samples of the dataset. The final prediction is the average (for regression) or the majority vote (for classification) of all trees, reducing the risk of overfitting. Random Forest is particularly effective for handling categorical variables and provides reliable estimates of feature importance (Equation 1).

$$\hat{y}_i = \frac{1}{K} \sum_{k=1}^K f_k(\mathbf{x}_i) \quad (1)$$

\hat{y}_i represents the predicted value for observation i .

K is the total number of trees in the ensemble of random forest.

f_k denotes the function of the k -th tree.

X_i is the feature vector for observation i .

2. EXtreme Gradient Boosting (XGBoost): XGBoost is an optimized implementation of gradient boosting that builds decision trees sequentially, minimizing a specified loss function at each step. It incorporates regularization techniques to prevent overfitting and is particularly effective at capturing non-linear relationships and handling imbalanced datasets [28] (Equation 2).

$$\hat{y}_i = \sum_{k=1}^K f_k(\mathbf{x}_i), \quad f_k \in \mathcal{F} \quad (2)$$

\hat{y}_i represents the predicted value for observation i .

K is the total number of trees in the ensemble.

f_k denotes the function of the k -th tree.

X_i is the feature vector for observation i .

\mathcal{F} is the set of learned decision trees.

3. Support Vector Machine (SVM): SVM aims to find the optimal hyperplane that separates data points of different classes while maximizing the margin between them. By employing kernel functions, SVM can map the data into higher-dimensional spaces to handle non-linear separability effectively (Equation 3).

$$\hat{y}_i = \text{sign} \left(\sum_{j=1}^n \alpha_j y_j K(\mathbf{x}_i, \mathbf{x}_j) + b \right) \quad (3)$$

\hat{y}_i represents the predicted value for observation i .

α_j are the coefficients associated with the support vectors.

y_j is the class label for observation j .

$K(\mathbf{x}_i, \mathbf{x}_j)$ is the kernel function evaluated between observations i and j .

b is the bias term.

n is the total number of support vectors.

2.3 Proposed approach for predictive modeling and risk assessment

Following the feature selection phase described in the previous section, the extracted feature sets were used across the selected machine learning models. The dataset was partitioned into three subsets: 20% allocated to the test set and 80% to the training-validation set. Following an adjustment, the training-validation subset was further split, with 80% used for training and 20% for validation, ensuring that model performance was validated on unseen data. To enhance the model results, Bayesian optimization was employed using the Optuna library [30], which efficiently explores the hyperparameter space to identify optimal combinations for each model (see Table 3). Additionally, an early stopping mechanism with a parameter "early_stopping_rounds=15" was implemented for all models, halting training if no improvement was observed after 15 iterations, thus preventing overfitting. The evaluation phase involved assessing model performance using four key metrics: precision (Equation 4), recall (Equation 5), F1-score (Equation 6), and the area under the ROC curve (AUC) (Equation 7).

$$\text{Precision} = \frac{TP}{TP + FP} \quad (4)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (5)$$

$$F1 - score = \frac{2 \times Precision \times Recall}{Precision + Recall} \quad (6)$$

$$AUC = \int_0^1 TPR(fpr)dfpr \quad (7)$$

To assess the impact of meteorological risk features from Météo-France and those derived from the large language model (LLM), the methodology was executed in two distinct stages. Initially, the analysis included F1 Space: General features combined with risk levels from Météo-France, followed by F2 Space: General features combined with risk levels derived from the LLM. This two-step approach facilitated a comprehensive comparison of the predictive capabilities of the different feature sets in forecasting firefighter intervention activity.

Table 3: Settings hyperparameters

Hyperparameter	XGBoost	Random forest	Support vector machine
max_depth	[2, 17]	[2, 17]	–
min_child_weight	[1, 100]	–	–
n_estimators	[100, 10000]	[100, 10000]	–
colsample_bytree	[0.3, 1.0]	–	–
learning_rate	[0.001, 0.3]	–	–
criterion	–	["gini","entropy"]	–
kernel	–	–	["rbf","sigmoid"]
C (regularization)	–	–	[0.01, 10]
gamma (kernel)	–	–	[0.001, 1.0]
class_weight	–	["balanced", None]	["balanced", None]

3 Results

In this study, the collection, processing, and compilation of a comprehensive dataset from various past intervention sources was time-consuming and required substantial effort. Particular attention was focused on meteorological risks, which were either derived from data provided by Météo-France or generated using the LLM, to ensure realistic and reliable results.

Following the training of the selected models and evaluation of their results, the outcomes were presented from various perspectives, ready for in-depth analysis in subsequent sections. Table 4 highlights the predictive performances of each model, focusing on the metrics described in the previous section, for the classification of firefighting activities based on the two feature spaces under study [F1, F2].

All models trained with F2 feature space showed reduced prediction errors and overall better performance compared to those trained with F1. Figures 5

Table 4: Prediction results for all models based on feature spaces F1 and F2

Model	Input	Precision	Recall	F1-score	AUC
XGBoost	F1	0.74	0.76	0.72	0.85
	F2	0.74	0.75	0.79	0.91
Random forest	F1	0.73	0.72	0.69	0.82
	F2	0.75	0.73	0.76	0.89
Support vector machine	F1	0.70	0.71	0.67	0.80
	F2	0.71	0.71	0.74	0.86

and 6 respectively illustrate the monthly mean prediction errors and the annual F1-score performance for the various models trained with either F1 or F2 feature spaces. Additionally, figure 7 provides clear evidence on the AUC performance that model XGBoost trained on F2 feature space outperformed those trained on F1 in predicting high and very high firefighter activity levels. However, for medium activity levels, the performance showed no significant differences, and both feature spaces produced similar results for low activity levels.

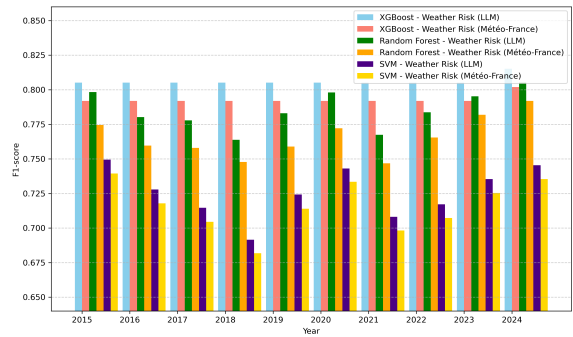


Fig. 5: F1-score performance by year for all models (F1 vs. F2)

Finally, Table 5 presents the optimized hyperparameters for each model, obtained through the Optuna framework. These parameters were crucial in enabling each model to achieve its best predictive performance.

4 Discussion

This study aimed to evaluate meteorological risk features derived either from Météo-France or generated using a Large Language Model (LLM) to improve predictions of firefighting interventions over the nine-year period (2015–2024). To achieve this, a detailed methodology was used, involving thorough data preparation, advanced analysis for informed decision-making, and feature selection.




Fig. 6: Average prediction errors by month for all models (F1 vs. F2)



Fig. 7: AUC performance by activity over the years

Table 5: Best hyperparameters

Hyperparameter	XGBoost [F1,F2]	Random forest [F1,F2]	Support vector machine [F1,F2]
max_depth	[5,2]	[17,15]	-
min_child_weight	[10,9]	-	-
n_estimators	1000	[500,490]	-
colsample_bytree	[0.8,0.7]	-	-
learning_rate	0.1	-	-
criterion	-	"gini"	-
kernel	-	-	"rbf"
C (regularization)	-	-	[1,2]
gamma (kernel)	-	-	[0.01,0.03]
class_weight	-	"balanced"	"balanced"

Initially, a general set of features was identified using various statistical and machine learning techniques. Subsequently, a second subset of features, focused on meteorological risks either derived from Météo-France or generated by LLM, was

extracted, resulting in two feature spaces: F1 and F2. To ensure meaningful and valuable results, three machine learning models (XGBoost, Random Forest, and Support Vector Machine (SVM)) were implemented to evaluate the predictive value of meteorological risk features.

The results demonstrated that models trained on feature space F2, which includes general features combined with meteorological risk features generated by LLM, outperformed others in terms of accuracy and reliability. In particular, XGBoost achieved the highest performance, with an F1-score of 0.79 and an AUC of 0.91 (see Table 4). Furthermore, the performance of this model showed consistent improvement year over year (see Figure 5). This trend suggests that meteorological risks generated by LLM may capture nuanced patterns, likely due to the model’s text-based analysis, which is capable of interpreting alarmist or descriptive information. However, this advantage could also reflect greater variability stemming from linguistic subtleties and the absence of direct calibration. A deeper analysis of monthly prediction errors revealed that, although models trained on feature space F2 generally performed better, they underperformed during the summer months compared to those using F1 (see Figure 6). This finding suggests that operational teams should focus on season-specific feature analysis in future studies. This discrepancy could arise due to limitations in the LLM’s training data, which may not comprehensively capture seasonal dynamics or low-activity scenarios. Furthermore, the reliance on descriptive prompts may amplify this gap, emphasizing the need for seasonal calibration and prompt refinement. Additionally, F2 inclusion significantly improved predictions for high and very high activity levels. However, for low activity levels, there was little difference between the two feature spaces (see Figure 7). This observation indicates that the detailed textual nature of LLM-generated risk scores might provide reliable insights, particularly for higher-risk scenarios. In contrast, further exploration of weak activity scenarios is warranted.

It is essential to highlight that the LLM’s effectiveness depends on the quality of its training data and prompt design. While the study leveraged robust datasets, the validation mechanisms for LLM-generated features remain an area requiring further exploration. Ensuring reliability in critical contexts demands rigorous evaluation protocols, which were not exhaustively discussed in this work. Addressing these aspects would enhance the credibility and operational applicability of such models. From a practical perspective, incorporating cost analysis reveals additional benefits. LLM-based models, despite their computational demands, offer relatively low operational costs when compared to alternative high-resolution meteorological data sources. Moreover, their deployment could lead to more efficient resource allocation, ensuring faster response times and mitigating severe consequences, such as property damage or loss of life, especially in high-risk scenarios. This cost-effectiveness, coupled with improved urgency identification, makes them a valuable tool for firefighting operations. To make the study’s findings more tangible, examples of practical impact could include scenarios such as the faster containment of wildfires due to earlier resource deployment or the improved management of simultaneous incidents during peak

seasons. These cases underscore the operational significance of integrating LLM-generated features.

Feature analysis is crucial for enhancing operational efficiency in fire intervention planning. A reliable predictive model would enable fire services to better anticipate resource needs, optimize operational logistics, and reduce response times, ultimately improving emergency management, safeguarding lives, and minimizing property damage.

In conclusion, the findings highlight the added value of integrating LLM-generated features into feature analysis, emphasizing the need to tailor approaches to specific scenarios to minimize losses. However, to achieve greater precision and optimization, further research is required. This includes exploring other LLM-based solutions, such as integrating media data for comparative analysis and leveraging deep learning techniques to enhance outcomes.

5 Conclusion

This study provides an in-depth analysis to enhance feature selection with a novel approach, optimizing the prediction performance of firefighter intervention activities. It leverages generative AI to extract relevant and optimal meteorological risk information. Using a comprehensive nine-year dataset from SDIS 25 in the Doubs region, France, this research tackles the integration of LLMs into feature generation and evaluates their strengths from multiple perspectives. The findings indicate that models trained on the F2 feature space, which combines general features with meteorological risk scores generated by LLMs, showed superior performance, especially for high and very high intervention activities. However, these models were less effective during the summer season and for low-activity scenarios. This insight could help operational personnel adopt this approach selectively under specific conditions.

As part of ongoing research efforts, a deeper exploration of media-based data is planned to enhance LLM-extracted feature sets. This will involve investigating alternative prediction perspectives and solutions. The primary objective remains to enhance predictive capabilities and optimize resource allocation by extracting the most relevant features for efficient firefighting interventions.

This work has been achieved in the frame of the EIPHI Graduate school (contract "ANR-17-EURE-0002").

References

1. Raiaan, et al. *A review on large Language Models: Architectures, applications, taxonomies, open issues and challenges*, *IEEE Access*, 2024, Publisher: IEEE.
2. Raweh, Abeer A and Nassef, Mohammed and Badr, Amr, *A hybridized feature selection and extraction approach for enhancing cancer prediction based on DNA methylation*, *IEEE Access*, vol. 6, pp. 15212–15223, 2018, Publisher: IEEE.

3. Malekipirbazari, Milad and Aksakalli, Vural and Shafqat, Waleed and Eberhard, Andrew, *Performance comparison of feature selection and extraction methods with random instance selection*, *Expert Systems with Applications*, vol. 179, p. 115072, 2021, Publisher: Elsevier.
4. Kabongo, Salomon and D'Souza, Jennifer and Auer, Sören, *Effective context selection in llm-based leaderboard generation: An empirical study*, *International Conference on Applications of Natural Language to Information Systems*, pp. 150–160, 2024, Publisher: Springer.
5. Lin, Xinyu and Wang, Wenjie and Li, Yongqi and Yang, Shuo and Feng, Fuli and Wei, Yinwei and Chua, Tat-Seng, *Data-efficient Fine-tuning for LLM-based Recommendation*, *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 365–374, 2024.
6. Cerna, Selene and Guyeux, Christophe and Royer, Guillaume and Chevallier, Céline and Plumerel, Guillaume, *Predicting fire brigades operational breakdowns: A real case study*. In: *Mathematics*, vol. 8, pp. 1383, 2020, Publisher: MDPI.
7. Cerna, Selene and Guyeux, Christophe and Laiymani, David, *The usefulness of NLP techniques for predicting peaks in firefighter interventions due to rare events*. In: *Neural Computing and Applications*, vol. 34, pp. 10117–10132, 2022, Publisher: Springer.
8. Sirri, Guyeux, *Firefighter Intervention Predictive Modeling: Reliability Assessment*. In: *International Conference on Circuit, Systems and Communication*, IEEE , 2024.
9. Sirri, Guyeux, *Air Quality Impact on Firefighter Interventions: Factors Analysis*. In: *The 7th International Conference on Big Data and Internet of Things*, Springer, 2024.
10. Sirri, Guyeux, *Solar activity Impact on Firefighter Interventions: Factors Analysis*. In: *The 5th International Conference on Deep Learning and Applications*, Springer, 2024.
11. Sirri, Guyeux, *River Levels Affecting Firefighter Interventions: Factor Analysis*. In: *The 28th International Database engineered applications symposium*, Springer, 2024.
12. Sirri, Guyeux, *Factor Analysis of Weather Conditions Impact on Firefighter Interventions*. In: *The 11th International Conference on Future Data and Security Engineering*, Springer, 2024.
13. The Sentinel Network. <https://www.sentiweb.fr/?page=table>. Accessed on 30 November 2024.
14. Astral. <https://pypi.org/project/astral/0.5/>. Accessed on 30 November 2024.
15. Skyfield. <https://github.com/skyfielders/python-skyfield>. Accessed on 30 November 2024.
16. NASA. <https://lance.modaps.eosdis.nasa.gov/viirs/>. Accessed on 30 November 2024.
17. MODIS. <https://lance.modaps.eosdis.nasa.gov/modis/>. Accessed on 30 November 2024.
18. Ministry of National Education. <http://www.education.gouv.fr/pid25058/le-calendrier-scolaire.html>. Accessed on 30 November 2024.
19. NASA. <https://www.swpc.noaa.gov/>. Accessed on 30 November 2024.
20. ATMO-BFC. <https://www.atmo-bfc.org/accueil>. Accessed on 30 November 2024.
21. Vigilance-France. <https://vigilance.meteofrance.fr/fr>. Accessed on 30 November 2024.
22. Pedregosa, et al. *Scikit-learn: Machine Learning in Python*. *Journal of Machine Learning Research*, 12, 2011, pp. 2825–2830.

23. Target Encoder. <https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.TargetEncoder.html>. Accessed on 30 November 2024.
24. Garreta, Raúl and Moncecchi, Guillermo. *Learning scikit-learn: machine learning in Python: experience the benefits of machine learning techniques by applying them to real-world problems using Python and the open source scikit-learn library*. (No Title), 2013.
25. Zien, Alexander and Krämer, Nicole and Sonnenburg, Sören and Rätsch, Gunnar. *The feature importance ranking measure*. In: *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2009, Bled, Slovenia, September 7-11, 2009, Proceedings, Part II*, pp. 694–709, 2009, Springer.
26. Ke, Guolin and Meng, Qi and Finley, Thomas and Wang, Taifeng and Chen, Wei and Ma, Weidong and Ye, Qiwei and Liu, Tie-Yan. *Lightgbm: A highly efficient gradient boosting decision tree*. In: *Advances in neural information processing systems*, volume 30, 2017.
27. Breiman, Leo. *Random forests*. In: *Machine learning*, volume 45, pp. 5–32, 2001, Publisher:Springer.
28. Chen, Tianqi and Guestrin, Carlos. *Xgboost: A scalable tree boosting system*. In: *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794, 2016.
29. Cortes, Corinna. *Support-Vector Networks*. In: *Machine Learning*, 1995.
30. Optuna. <https://optuna.org/>. Accessed on 30 November 2024.