

# Predicting the Category of Fire Department Operations

Kevin Pirklbauer

kevin.pirklbauer@students.fh-hagenberg.at  
Department of Mobility & Energy, University of Applied  
Sciences Upper Austria  
Hagenberg, Austria

Rainhard Dieter Findling

rainhard.findling@aalto.fi  
Ambient Intelligence Group, Department of  
Communications and Networking, Aalto University  
Espoo, Finland

## ABSTRACT

Voluntary fire departments have limited human and material resources. Machine learning aided prediction of fire department operation details can benefit their resource planning and distribution. While there is previous work on predicting certain aspects of operations within a given operation category, operation categories themselves have not been predicted yet. In this paper we propose an approach to fire department operation category prediction based on location, time, and weather information, and compare the performance of multiple machine learning models with cross validation. To evaluate our approach, we use two years of fire department data from Upper Austria, featuring 16.827 individual operations, and predict its major three operation categories. Preliminary results indicate a prediction accuracy of 61%. While this performance is already noticeably better than uninformed prediction (34% accuracy), we intend to further reduce the prediction error utilizing more sophisticated features and models.

## CCS CONCEPTS

• **Human-centered computing** → **Ubiquitous and mobile computing**; *Visualization*; • **Computing methodologies** → *Machine learning*.

## KEYWORDS

machine learning, fire department operation prediction, operation category prediction

### ACM Reference Format:

Kevin Pirklbauer and Rainhard Dieter Findling. 2019. Predicting the Category of Fire Department Operations. In *The 21st International Conference on Information Integration and Web-based Applications & Services (iiWAS2019)*, December 2–4, 2019, Munich, Germany. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3366030.3366113>

## 1 INTRODUCTION

Around 20.000 operations have been carried out by voluntary fire departments in Upper Austria in 2017 and 2018. Since both human and material resources of voluntary fire departments are limited, distribution of their resources is critical. Machine learning has

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*iiWAS2019, December 2–4, 2019, Munich, Germany*

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 978-1-4503-7179-7/19/12...\$15.00  
<https://doi.org/10.1145/3366030.3366113>

been used to assist fire departments by predicting the targets the departments are concerned with, such as forest fires [2, 25, 28], the number of fire trucks needed for operations [8, 9], the risk of building fires happening [19, 20], or the chances a traffic incident occurs [10, 24, 33].

To further assist fire departments in their resource planning we extend existing prediction approaches by adding a fire department operation category prediction. Given that an operation happens or is predicted to happen, then our system can predict the corresponding operation category. This allows for more precise resource planning and allocation through prediction cascades: a resource requirement prediction model could use the operation category probability as feature. Alternatively, resource prediction could be done by individual models, where each predicts resources for exactly the assessed operation category.

In this paper we propose an approach to predict the three major categories of fire department operation categories in Upper Austria. In our evaluation we utilize a dataset of fire department operations from 2017-2018, where each operation contains multiple features, on which we train and compare multiple machine learning models with cross-validation. Summarizing, the contributions in this paper are:

- We propose an approach to predict the three major categories of fire department operation categories in Upper Austria based on time, location, and weather data.
- We give insight into the correlation between weather data features and fire department operation categories.
- We evaluate our approach with a dataset of 16.827 fire department operations from 2017-2018, and compare the performance of multiple machine learning models with cross-validation.

## 2 RELATED WORK

Previous work on machine learning aided prediction for fire department related operations has mainly investigated fire and traffic incident prediction, as well as incident duration prediction. In this section we highlight previous approaches in those three areas that are related to or important for our work. Previous work has also investigated further areas that are less strongly related to fire department operations, which we do not discuss in detail in this section. Examples include general emergency incident prediction [22, 26] or thunderstorm incident prediction [14, 16, 23].

### 2.1 Fire incident prediction

Fire incident prediction has mostly focused on the amount of fire trucks required for an incident, as well as on predicting risk of wildfire and building fires. The number of fire trucks required per

region and firestation on one day in Amsterdam-Amstelland is predicted in [8]. The authors use historical fire department operation data combined with weather features, like wind speed, temperature, rainfall, and visibility. They evaluate four different regression models where an ensemble averaging (EA) model performs best with a mean absolute percentage error (MAPE) of 0.1853. A similar approach is presented in [9], in which the authors predict the number of firetrucks needed by vehicle type based on demographic features, like population density, in addition to weather features. In their evaluation a gradient boosting model performs best with a mean squared error (MSE) of 0.0953. Firebird [19] is a framework for predicting building fire risk in Atlanta based on fire incident data, property fire risk inspection data, and structural data of properties (e.g. age, number of floors, and condition). A random forest (RF) model performs best with an area under curve (AUC) value of 0.8246. Building upon the Firebird approach, [20] propose another framework for building fire risk prediction in Pittsburgh, Pennsylvania. In their evaluation an XGBoost model achieves a Kappa value of 0.37. [28] uses remote sensing in combination with a support vector machine (SVM) and artificial neural network (ANN) model to predict wildfires, achieving an accuracy of 97.84% for the SVM and 98.32% for the ANN. [25] map wildfires to a 10×10 km grid and use SVM models to predict wildfire risk for each tile on the grid based on environmental factors, achieving an accuracy of 96%. Other wildfire prediction approaches include [27], [7] and [6], using SVM, RF, and genetic programming models.

## 2.2 Traffic incident prediction

In traffic accident prediction, many approaches focus on predicting if accidents will happen, or on the amount of accidents that will happen. A popular approach for data collection for traffic incident prediction are loop detectors. Loop detectors use induction loops embedded in the ground to detect vehicles entering or passing through an area (e.g. a freeway segment or a traffic light). The prediction models used for these approaches are SVM [17, 32], ANN [32], RF [1], and k-Nearest-Neighbor (kNN) [18] models. [4] predict traffic incidents in the state of Iowa based on traffic data, road information, and weather data using a RF model. They thereby achieve an accuracy of 65%. Similarly, the approach in [33] predicts traffic incidents in Iowa based on rainfall data, road network information, and demographic data. It achieves an AUC of 0.9612 with an ANN and SVM model. [10] proposes a hybrid approach that combines unsupervised and supervised deep learning with traffic accident data, road pavement factors, and environmental characteristics to predict traffic incidents in Tennessee. It achieves a mean absolute error (MAE) of 0.660. Traffic incidents are mapped to a 1×1 km grid in [24] to predict incident risks for each tile in the grid. Their recurrent neural network (RNN) model achieves a root mean squared error (RMSE) of 0.034.

## 2.3 Incident duration prediction

To predict traffic incident duration, [31] uses information about physical traits of incidents, response and traffic measurements taken, and the types of incidents. Incidents comprise the groups "stopped vehicle", "lost load", and "accident" since important variables for predicting the target seem to differ between those groups.

Prediction is done by different partial least squares regression models with the best model achieving an accuracy of 71.30%. ANN models are used in [11] to predict traffic incident duration on a freeway in Guangzhou, China. The authors do not report the regression error directly, but a low correlation of 0.8535 between their training- and test-set prediction target values. Based on this they argue that the amount of randomness in their traffic incident data should disallow improving upon their results. Other approaches to incident duration prediction include [3], which proposes the use of hazard-based models for incident duration prediction (MAPE 43.7%) and compares their performance to kNN models (MAPE 41.1%). Furthermore, [5] uses Naïve Bayesian classifiers for incident duration prediction (56% accuracy) and compares the results to a linear regression model (45% accuracy). More thorough comparisons between multiple incident duration prediction approaches and different models are given in [30] and [15].

Summarizing, there is numerous previous work to predict if traffic or fire incidents will happen, to predict the amount of such incidents, or the resources required for an incident. However, research on predicting fire department operation categories seems sparse. We are not aware of any approach that – given that an incident happens or is predicted to happen – would predict the corresponding fire department operation category. We argue that such an approach would allow for subsequently more precisely predicting the required resources for the incident, and consequently aid fire departments in their resource planning and allocation.

## 3 DATASET DESCRIPTION

Our dataset has been collected via the publicly available API<sup>1</sup> of the Oberösterreichischer Landes-Feuerwehrverband (OOELFV), which is the fire department umbrella organization in Upper Austria. The API allows for querying fire department operations for past three months. The complete dataset has been collected over nearly 2.5 years, from 10/2016 to 02/2019, and contains 25.136 fire department operations in Upper Austria. To include data from only fully covered year cycles in our analysis, we chose to limit our data to the 21.419 samples that are within the years 2017 and 2018. Each fire department operation in that dataset contains information about the location, start and end time and date, alarm level, as well as category and type of the operation. The category is one of five primary categories (German original categories in brackets): fire (Brand), technical (Technisch), storm (Unwetter), person (Person), misc. (Sonstige). The type is one of 9 secondary types and one of 46 secondary subtypes, which specify more fine granular categories for the operation. Within those categories and types we focus on predicting the primary category of an operation. Those are imbalanced, with fire operations making up about 42% of all samples, while 27% and 14% of operations belong to the categories technical and storm (Fig. 1(a)). Operations seem to be more frequent in the morning and evening hours (Fig. 1(c)), on Mondays and Fridays (Fig. 1(b)), as well as in summer and winter (Fig. 1(d)).

## 4 APPROACH

Our approach to predict the fire department operation category consists of preprocessing, including adding weather data, feature

<sup>1</sup>OOELFV intranet: <https://intranet.oelfv.at/>.

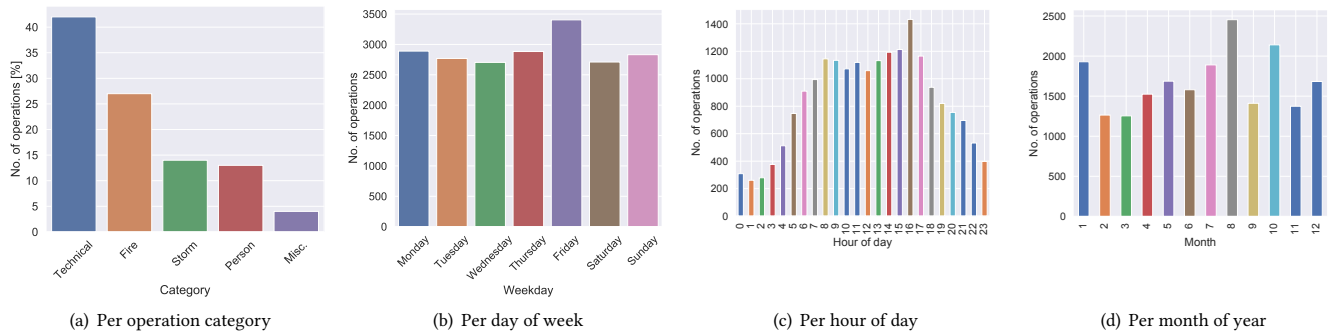


Figure 1: Fire department operation frequencies over different features in our dataset.



Figure 2: Overview of the processing chain employed in our fire department operation category prediction.

selection, adding further features, and data normalization, as well as dimensionality reduction and a model tuning with hyperparameter grid-search (Fig. 2). In this section we describe the details of those steps. The analysis and evaluation is done with Python-based Jupyter [12] and Scikit-learn [21] for model training and pre-processing.

### 4.1 Weather Data Addition

For our approach we focus on the three major fire department operation categories in the dataset, namely fire, technical, and storm, which have a combined total of 18.055 operation samples. For each sample we try to acquire weather data via DarkSky<sup>2</sup>, based on its start time and geolocation. We add the following weather information: precipitation probability, precipitation intensity, temperature, apparent temperature, dew point, humidity, pressure, wind speed, wind bearing, wind gust, cloud cover, UV index, and visibility. Weather data was fully unavailable for 1.228 samples (6.8% of the dataset), which we exclude from our analysis as a consequence. The remaining 16.827 operation samples form the basis for our evaluation. For parts of those samples, weather information was only partially available. We use kNN feature imputation [29] to approximate the missing features for those samples (number of imputed samples in brackets): precipitation intensity and probability (both 5458), pressure (5462), wind speed and bearing (both 294), wind gust (5324), cloud cover (2440), UV index (1529), and visibility (2461).

### 4.2 Feature Selection

To get insight to which features might be important for our classification task we perform data visualization and a visual feature analysis. We thereby observe a slight difference in the spread of samples of our categorical target variable with multiple features. Those include, for example, pressure and dew point, where storm operation frequency seems to be higher with lower air pressure

<sup>2</sup>DarkSky: <https://darksky.net/dev>.

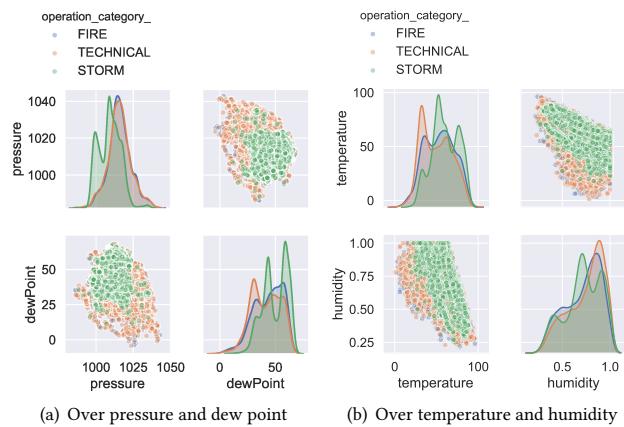


Figure 3: Distribution of operation category in weather data feature space.

and higher dew point (Fig. 3(a)). They also include temperature and humidity, where the share of technical operations is higher with lower temperature (Fig. 3(b)). While the visible differences in the distribution patterns seem small, they are supported by intuition, and the accumulated differences over all feature space dimensions might indicate predictive power for our target variable.

On top of the obtained features we apply four different feature selection methods [13]. Those include selectKbest, which removes all but the best  $k$  features, measured by their ANOVA- $f$  values, as well as wrapper methods [13], which determine feature importance based on model coefficients. Additionally we train a decision tree and random forest model with our data, then use the resulting model structure to determine the corresponding feature importance. After comparing the results of the four feature selection approaches, we decide to use the 8 most important numerical weather features:

temperature in degree Celsius and Kelvin, the apparent temperature in Kelvin, pressure, dew point, wind speed, wind gust, and precipitation intensity.

### 4.3 Feature Addition

Based on the selected 8 features we add further feature representations to the dataset. Those include their squared, logarithmic, and square-rooted forms as suggested in [13]. Furthermore, since our features also comprise cyclic features like hour-of-day, month-of-year, or day-of-week, we also add a sine and cosine representation of those to the dataset. We thereby map one full feature cycle to orthogonal numeric sine and cosine representations with the sin and cos trigonometric functions in their corresponding unit circle. At the example of hour-of-day, midnight would map to a sine and cosine representation of (0,1), 6 am to a representation of (1,0), 12 am to (0,-1), and 6 pm to (-1,0). This transformation preserves the underlying information of the cyclic feature and represents them in two numeric values, while avoiding a huge value change when one cycle ends and the next cycle begins, e.g. hour 23 to hour 0 in the hour-of-day example.

### 4.4 Feature Normalization and Reduction


For normalization we use a simple center-scale method that normalizes features to mean = 0 and  $\sigma = 1$ . We further utilize PCA for dimensionality reduction. The number of extracted PC components to be used is determined by the hyperparameter grid-search of the model evaluation, which in our evaluation was determined to be 49. Since many machine learning algorithms cannot handle categorical variables, like the operation category in our case, we encode it with a label encoder [13] that assigns an integer value  $[0, n - 1]$  to each of the  $n$  categories. After feature extraction and preprocessing, our dataset consists of 16.827 fire department operation samples, each comprising a total of 49 features.

### 4.5 Recognition and Results

We use the resulting 16.827 fire department operation samples, each containing 49 features, to train and evaluate our operation category prediction. We utilize and compare different machine learning algorithms, including kNN, Linear Discriminant Analysis (LDA), DT, RF, and SVM with an Radial Basis Function (RBF) kernel. We compare those to multiple uninformed baseline models, including a stratified model (predictions based on target class distribution), a most-frequent model (always predicts the most frequently occurring class), a uniform model (predicts uniformly random), and a prior model (predictions based on prior distribution of classes). The baseline model that performs best is the stratified model, which in cross-validation achieves a mean prediction performance of 34% accuracy. In comparison, our models achieve noticeably better cross-validation results (Fig. 4), where the SVM and RF models perform nearly twice as well, both with a mean accuracy of 61% ( $\sigma = 2.2\%$ ).

### 4.6 Discussion

While the models evaluated in our approach noticeably outperform the corresponding baseline models, overall performance still seems improvable. The current prediction errors could be due to several reasons. First, it might be that the employed features and models are



**Figure 4: Grid-search cross-validation results for all models. Error bars depict the standard deviation.**

not yet powerful enough. That our models outperform uninformed prediction indicates that the data contains information about our prediction target, the operation category. Still, we deem it likely that the information contained in the current features still is limited. In this case, the performance will likely improve by adding further information with relation to our prediction target to the dataset, like traffic data features, demographic features, or by transforming the contained information with further and more sophisticated feature engineering. It might also be that the relations between features and prediction target are too complex for the currently employed modes. In this case, employing more complex models, like Deep Neural Networks, might also improve the prediction performance. We assume that the current result is caused less by this than by the previous point. Second, data imputation for missing features might also have negative influence. This could be improved by utilizing a complete dataset without missing values, which might be available from official government sources like the Zentralanstalt für Meteorologie und Geodynamik (ZAMG) in Austria.

To summarize, we deem the current performance of our approach to be too low to be employed for fire department operation category prediction in production environments. However, our first results seem promising, and we argue that our approach serves as basis for further research towards this goal. Moreover, we are confident that by increasing the information contained in the dataset (both in features as well as in completeness), and by employing more sophisticated prediction models, future research can improve results towards an applicable fire department category prediction.

## 5 CONCLUSION

In this paper we have presented an approach to predict fire department operation categories in Upper Austria, based on location, time, and weather data. Our processing chain comprises data and feature addition, data normalization, dimensionality reduction with PCA, as well as model training and selection with a hyper-parameter grid-search and cross-validation. For our evaluation we utilize two years of fire department operations data from Upper Austria. After processing, we use 16.827 operation samples with 49 features each in our evaluation. There we consider multiple machine learning models, including kNN, LDA, DT, RF, and SVM, and compare their prediction performance with multiple uninformed baseline models.

Data visualization shows relations between weather data and operation categories, indicating potential in their predictive power. Our models outperform the baseline models. Our SVM and RF models achieve an accuracy of 61%, which is nearly twice the performance of the best baseline model (34% accuracy). Still, we deem the current performance to be insufficient for reliable operation category prediction in a production environment. We attribute the current prediction performance to multiple factors and shortcomings, including improvable amount of information contained in the dataset, insufficiently strong prediction models, as well as information loss due to data imputation for missing values. However, we also argue that those first results seem promising for improving the fire department operation category prediction in future work, among others by addressing the mentioned shortcomings. For future work based on our approach and results we would suggest more sophisticated features in the dataset as well as feature engineering, including, for example, traffic data, demographic data, and alike. We further suggest to consider stronger machine learning models, such as DNNs, to model the more complex relationships of those features for prediction.

## ACKNOWLEDGMENTS

We appreciate funding in the frame of the ChistEra RadioSense Big Data and process modeling for Smart Industry (BDSI) project.

## REFERENCES

- [1] Mohamed M. Ahmed and Mohamed A. Abdel-Aty. 2011. The Viability of Using Automatic Vehicle Identification Data for Real-Time Crash Prediction. *IEEE Transactions on Intelligent Transportation Systems* 13, 2 (June 2011), 459–468. <https://doi.org/10.1109/ITIS.2011.2171052>
- [2] Amparo Alonso-Betanzos, Oscar Fontenla-Romero, Bertha Guijarro-Berdinas, Elena Hernandez-Pereira, Maria Inmaculada Paz Andrade, Eulogio Jimenez, Jose Luis Legido Soto, and Tarsy Carballas. 2003. An intelligent system for forest fire risk prediction and fire fighting management in Galicia. *Expert Systems with Applications* 25, 4 (2003), 545–554. [https://doi.org/10.1016/S0957-4174\(03\)00095-2](https://doi.org/10.1016/S0957-4174(03)00095-2)
- [3] Bahar N. Araghi, Simon Hu, Rajesh Krishnan, Michael Bell, and Washington Ochieng. 2014. A comparative study of k-NN and hazard-based models for incident duration prediction. In *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)*. 1608–1613. <https://doi.org/10.1109/ITSC.2014.6957923>
- [4] Mohana Bansal, Kushal Thakkar, and Gaurav Shahane. 2017. *Prediction Model of Potential Road Crashes Bases on the Spatial and Temporal Factors*. Technical Report. University of Maryland College Park.
- [5] Stephen Boyles, David Fajardo, and S. Travis Waller. 2015. Naive Bayesian classifier for incident duration prediction.
- [6] Mauro Castelli, Leonardo Vanneschi, and Aleš Popović. 2015. Predicting Burned Areas of Forest Fires: an Artificial Intelligence Approach. *Fire Ecology* 11, 1 (01 Apr 2015), 106–118. <https://doi.org/10.4996/fireecology.1101106>
- [7] Paulo Cortez and Anibal Morais. 2007. A Data Mining Approach to Predict Forest Fires using Meteorological Data. *DSI - Engenharia da Programação e dos Sistemas Informáticos* (Dec. 2007).
- [8] Jeffrey de Deijn. 2017. *Predicting fire brigade incidents*. Brandweer Amsterdam-Amstelland, Karspeldreef 16 1090 AD Amsterdam.
- [9] Omesh Debipersad. 2017. *Predicting the use of fire brigade materials at an incident*. Brandweer Amsterdam-Amstelland, Karspeldreef 16 1090 AD Amsterdam.
- [10] Chunjiao Dong, Chunfu Shao, Juan Li, and Zhihua Xiong. 2018. An Improved Deep Learning Model for Traffic Crash Prediction. *Journal of Advanced Transportation* 2018 (Dec. 2018), 1–13. <https://doi.org/10.1155/2018/3869106>
- [11] Liping Guan, Weiming Liu, Xiangyan Yin, and Luping Zhang. 2010. Traffic Incident Duration Prediction Based on Artificial Neural Network. In *2010 International Conference on Intelligent Computation Technology and Automation*, Vol. 3. 1076–1079. <https://doi.org/10.1109/ICICTA.2010.418>
- [12] Thomas Kluyver, Benjamin Ragan-Kelley, Fernando Pérez, Brian Granger, Matthias Bussonnier, Jonathan Frederic, Kyle Kelley, Jessica Hamrick, Jason Grout, Sylvain Corlay, Paul Ivanov, Damián Avila, Safia Abdalla, and Carol Willing. 2016. Jupyter Notebooks - A publishing format for reproducible computational workflows. In *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, F. Loizides and B. Schmidt (Eds.). IOS Press, 87–90. <https://doi.org/10.3233/978-1-61499-649-1-87>
- [13] Max Kuhn and Kjell Johnson. 2013. *Applied Predictive Modeling*. Springer New York. <https://books.google.at/books?id=xYRDAAAQBAJ>
- [14] Ryan Lagerquist, Cameron Homeyer, Amy Mcgovern, Corey Potvin, Thea Sandmael, and Travis Smith. 2019. Development and interpretation of deep-learning models for nowcasting convective hazards. <https://doi.org/10.13140/RG.2.2.22218.82885>
- [15] Ruimin Li, Francisco C. Pereira, and Moshe E. Ben-Akiva. 2018. Overview of traffic incident duration analysis and prediction. *European Transport Research Review* 10, 2 (may 2018). <https://doi.org/10.1186/s12544-018-0300-1>
- [16] A. J. Litta, Sumam Idicula, and C. Naveen Francis. 2012. Artificial Neural Network Model for the Prediction of Thunderstorms over Kolkata. *International Journal of Computer Applications* 50 (07 2012), 50–55. <https://doi.org/10.5120/7819-1135>
- [17] Yisheng Lv, Shuming Tang, Hongxia Zhao, and Shuang Li. 2009. Real-time highway accident prediction based on support vector machines. In *2009 Chinese Control and Decision Conference*. 4403–4407. <https://doi.org/10.1109/CCDC.2009.5192409>
- [18] Yisheng Lv, Shuming Tang, Hongxia Zhao, and Shuang Li. 2009. Real-Time Highway Traffic Accident Prediction Based on the k-Nearest Neighbor Method. In *2009 International Conference on Measuring Technology and Mechatronics Automation*, Vol. 3. 547–550. <https://doi.org/10.1109/ICMTMA.2009.657>
- [19] Michael Madaio, Shang-Tse Chen, Oliver L. Haimson, Wenwen Zhang, Xiang Cheng, Matthew Hinds-Aldrich, Duen Horng Chau, and Bistra Dilkina. 2016. Firebird: Predicting fire risk and prioritizing fire inspections in Atlanta. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*. ACM, New York, NY, USA, 185–194. <https://doi.org/10.1145/2939672.2939682>
- [20] Michael A. Madaio. 2018. *Predictive Modeling of Building Fire Risk*. Technical Report. Metro21 Research Publication.
- [21] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [22] Geoffrey Pette, SaideepNannapaneni, Benjamin Stadnick, Abhishek Dubey, and Gautam Biswas. 2017. Incident analysis and prediction using clustering and Bayesian network. In *2017 IEEE SmartWorld, Ubiquitous Intelligence Computing, Advanced Trusted Computed, Scalable Computing Communications, Cloud Big Data Computing, Internet of People and Smart City Innovation*. 1–8. <https://doi.org/10.1109/UIC-ATC.2017.8397587>
- [23] Sundaramoorthy Rajasekaran, S. Gayathri, and T.-L. Lee. 2008. Support vector regression methodology for storm surge predictions. *Ocean Engineering* 35, 16 (nov 2008), 1578–1587. <https://doi.org/10.1016/j.oceaneng.2008.08.004>
- [24] Honglei Ren, You Song, JingXin Liu, Yucheng Hu, and Jinzhi Lei. 2017. A Deep Learning Approach to the Prediction of Short-term Traffic Accident Risk. *CoRR* abs/1710.09543 (2017). [arXiv:1710.09543](https://arxiv.org/abs/1710.09543)
- [25] Marcos Rodrigues and Juan de la Riva. 2014. An insight into machine-learning algorithms to model human-caused wildfire occurrence. *Environmental Modelling & Software* 57 (jul 2014), 192–201. <https://doi.org/10.1016/j.envsoft.2014.03.003>
- [26] Tyler Romero, Zachary Barnes, and Frank Cipollone. 2017. *Predicting emergency incidents in San Diego*. Technical Report. CS229. Stanford School of Engineering.
- [27] George E. Sakr, Imad H. Elhadj, George Mitri, and Uche Wejinya. 2010. Artificial intelligence for forest fire prediction. In *2010 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*. 1311–1316. <https://doi.org/10.1109/AIM.2010.5695809>
- [28] Younes Oulad Sayad, Hajar Mousannif, and Hassan Al Moatassime. 2019. Predictive modeling of wildfires: A new dataset and machine learning approach. *Fire Safety Journal* 104 (mar 2019), 130–146. <https://doi.org/10.1016/j.firesaf.2019.01.006>
- [29] Volker Tresp, Ralph Neuneier, and Subutai Ahmad. 1995. Efficient Methods for Dealing with Missing Data in Supervised Learning. In *Advances in Neural Information Processing Systems 7*, G. Tesauro, D. S. Touretzky, and T. K. Leen (Eds.). MIT Press, 689–696.
- [30] Gaetano Valentini, Maria Lelli, and Domenico Cucina. 2010. A comparative study of models for the incident duration prediction. *European Transport Research Review* 2, 2 (01 Jun 2010), 103–111. <https://doi.org/10.1007/s12544-010-0031-4>
- [31] Xuanqiang Wang, Shuyan Chen, and Wenchang Zheng. 2013. Traffic Incident Duration Prediction Based On Partial Least Squares Regression. *Procedia - Social and Behavioral Sciences* 96 (11 2013). <https://doi.org/10.1016/j.sbspro.2013.08.050>
- [32] Baozhen Yao, Ping Hu, Mingheng Zhang, and Maoqing Jin. 2014. A Support Vector Machine with the Tabu Search Algorithm for Freeway Incident Detection. *Int. J. Appl. Math. Comput. Sci.* 24, 2 (June 2014), 397–404. <https://doi.org/10.2478/amcs-2014-0030>
- [33] Zhenghong Yuan, Xun Zhou, Tianbao Yang, and James Tamerius. 2017. Predicting Traffic Accidents Through Heterogeneous Urban Data : A Case Study. In *6th International Workshop on Urban Computing*.