Lessons from past hazardous events: data analytics for severity prediction

Nicola Paltrinieri

Department of Mechanical and Industrial Engineering, Norwegian University of Science and Technology NTNU, S. P. Andersens veg 3, Trondheim, Norway. E-mail: nicola.paltrinieri@ntnu.no

Riccardo Patriarca

Department of Mechanical and Aerospace Engineering, Sapienza University of Rome, Via Eudossiana, 18, Rome, Italy. E-mail: riccardo.patriarca@uniroma1.it

Michael Pacevicius

Department of Mechanical and Industrial Engineering, Norwegian University of Science and Technology NTNU, S. P. Andersens veg 3, Trondheim, Norway. E-mail: michael.f.pacevicius@ntnu.no

Pierluigi Salvo Rossi

Department of Electronic Systems, Norwegian University of Science and Technology NTNU, O. S. Bragstads plass 2, Trondheim, Norway. E-mail: pierluigi.salvorossi@ntnu.no

An increasing amount of information is collected by the monitoring systems within the process industry, especially concerning safety management. For instance, the Seveso III regulation on the control of major-accident hazards involving dangerous substances is the first version that refers to the collection of safety indicators for monitoring the performance of safety management systems. This leads to a call for improvement in learning past lessons and definition of techniques to process relevant data, in order to deal with unexpected events and provide the right support to safety management. Through this work, we suggest a data analytics approach for severity prediction of future hazardous events. The approach is twofold and is based on the use and comparison of multiple linear regression (MLR) and deep neural network (DNN) models. These models are developed and tested on the Major Hazardous Incident Data Service (MHIDAS) database. A set of simulations has been carried out not only to evaluate the models, but also to identify their limitations. The results show the capability of these models to manage heterogeneous data from past accident records and extract important information to support safety-related decision making. It must also be mentioned that intrinsic model limitations should be considered, and appropriate model selection and customization should be carefully carried out to deliver the needed support.

Keywords: learning from lessons, data analytics, multiple linear regression, deep neural network, safety management.

1. Introduction

Hazardous events may manifest under various forms in industry, but they mostly known as events involving the loss of containment of hazardous materials in the process industry (Pasman, 2015). In Europe, the handling of hazardous materials by process industry is regulated by the so-called the Seveso III regulation on the control of major-accident hazards involving dangerous substances (European Parliament and Council, 2012). As stated by Pasman, the loss of control of such substances has the potential to cause high-impact low-probability accidents (Pasman, 2015). High impact indicates catastrophic losses, but, due to their low probability, these accidents may even not happen during a lifetime.

Paltrinieri et al. (Paltrinieri et al., 2013, 2011) address another aspect of major accidents, as, in some cases, they are the results of scenarios that are "not captured by hazard identification methodologies because deviating from normal expectations of unwanted events or worst case reference scenarios." These may occur when hazard identification does not produce a complete overview of system hazards (Paltrinieri et al., 2010).

Another term used to define rare catastrophic events that have never been encountered before was coined by Taleb (Taleb, 2007), who used the metaphor of the Black Swans. These events can be explained only in the aftermath and cannot be anticipated, such as the black swan was believed to be impossible before its discovery in the 17th century (Taleb, 2007). However, the concept may be misused as it may represent a reason for

Proceedings of the 30th European Safety and Reliability Conference and the 15th Probabilistic Safety Assessment and Management Conference
Edited by Piero Baraldi, Francesco Di Maio and Enrico Zio
Copyright © ESREL2020-PSAM15 Organizers.*Published by* Research Publishing, Singapore.
ISBN: 978-981-14-8593-0; doi:10.3850/978-981-14-8593-0_4794-cd

ignoring the potential for major accidents and avoiding the implementation of long-term safety measures (Paté-Cornell, 2012).

The concept of Dragon-Kings (Sornette, 2009), may indicate a responsible approach to deal with major accidents. Dragon-Kings are defined as events that are extreme and outliers (in analogy with the kings' wealth), but unlike anything else, such as dragons. These major accidents are intended as the result of some degrees of organization and coordination of relatively smaller unwanted events and features, which could serve to amplify the final consequences.

Extreme accidents are the result of a combination of such details, some of which may be considered as deviations from normal/optimal conditions.

These deviations can be defined as early warnings (Paltrinieri et al., 2015a) or associated with the concept of "Small Things" (Paltrinieri and Khan, 2016). Small things might be recurring old issues in a plant or organization, which do not need imaginative definitions to be prevented, but perhaps only the compliance with already present procedures. Acting on Small Things would allow breaking the chain of events and lower the probability for major unpredictable accidents.

In the last decade, increasing attention has been dedicated to evaluation and monitoring of early deviations through appropriate indicators, to assess and control risk. Indicators can be represented by a series of factors: physical conditions of a plant (equipment pressure and temperature); number failures of an equipment piece; maintenance backlog; number of emergency preparedness exercises; amount of overtime worked; etc. Several indicator typologies have been theorized and used, but we often address risk indicators if (Øien, 2001): they provide numerical values (such as a number or a ratio); they are updated at regular intervals; they only cover some selected determinants of overall risk, in order to have a manageable set of them.

The latter feature has quickly become outdated due to the extensive collection that is being carried out in industry and the attempts made to process and elaborate larger numbers of them. For instance, for the first time since the first Seveso directive was issued in 1982, Seveso III mentions specific procedures for safety performance indicators and/or other relevant indicators, to use for monitoring the performance of safety management systems [3].

Table 1 reports how such suggestion has been received in some of the EU member and associated countries. Past hazardous events are collected by all the countries considered in Table 1. United Kingdom and France use specific databases to collect them. On the other hand, the use of safety performance indicators is not as common across Europe, but where it is not present, it is suggested by relevant research institutes.

Table 1. Seveso III-based monitoring approaches in the EU member and associated countries (Paltrinieri and Reniers, 2017).

	Indicators	
	Past events	*Safety performance indicators*
United Kingdom	Hazardous events reported to the competent authorities and regulated by RIDDOR (Reporting of Injuries, Diseases and Dangerous Occurrences).	The British competent authorities require hazard establishments to collect safety performance indicators (PSPIs).
France	Hazardous events are collected in the database ARIA (Analysis, Research and Information on Accidents). These events are also used as Key Performance Indicators.	The French national competence centre for industrial safety and environmental protection (INERIS) suggests the use of a Safety Performance Indicator System.
Italy	Hazardous events are reported to the competent authority.	The regulation states that safety performance monitoring should be based on indicators.
Netherlands	Hazardous events are reported to and collected by the competent authority.	The regulation requires companies handling hazardous substances to collect safety performance indicators.
Finland	The competent authority uses hazardous events to assess the performance of safety management systems.	Other indicators are also used to assess the performance of safety management systems.
Norway	The competent authority collects hazardous events according to the Seveso regulations.	The Norwegian research institute SINTEF suggests monitoring the safety trend of Seveso establishments with safety performance indicators.

This leads to a call for improvement in learning past lessons and definition of techniques to process relevant data, in order to deal with unexpected events and provide the right support to safety management. However, industrial risk analysis is unevenly progressing within this topic (Paltrinieri et al., 2019). At the same time, the use of machine learning has possibly become more attractive, given the progressive refinement of its models and the exponential increase in available computing power (Goodfellow et al., 2016). For this reason, we suggest a data analytics approach to predict the severity of potential hazardous events based on lessons learned from past hazardous events.

2. Method

The approach is twofold and is based on the use and comparison of multiple linear regression (MLR) and deep neural network (DNN) models. MLR and DNN may be considered as techniques belonging to the field of machine learning, which refers to techniques aiming to program computers to learn from experience (Samuel, 1959). While MLR is a statistical technique that uses several explanatory variables to predict the outcome of a response variable (Andrews, 1974), DNN aims to simulate (to a certain extent) the learning model of the human brain (Goodfellow et al., 2016). It is loosely based on information processing and communication patterns in a neural system. It allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction.

A computer may be trained to assess risk or some of its components for safety-critical industries such as Seveso-regulated sites through machine learning techniques. This would allow processing the large amount of information currently collected in the form of lessons from past events or safety performance indicators. Moreover, although risk level cannot be evaluated with certainty, machine learning can allow for expert supervision through supervised learning (Goodfellow et al., 2016).

2.1 Multiple linear regression

A linear model, given a vector of inputs $X = (x_1, x_2, \dots, x_p)$, predicts the output y (in this case an index for the risk R) via the following equation (Hastie et al., 2009):

$$y = b_0 + \sum_{j=1}^{p} x_j w_j \sim R \qquad (1)$$

where b_0 is the so-called bias and wj represents the model weights. This model needs then to be trained with a training set data in order to learn the weights of every provided input. Once the weights are known, the model can be used for prediction of y based on new inputs X.

2.2 Deep neural network

The deep learning model considered in this work is a feed-forward neural network, wherein connections between the units do not form a cycle (Svozil et al., 1997). A linear model, such as MLR, would be restricted to linear functions, while a DNN model describes the target as a nonlinear function of the input features (Goodfellow et al., 2016). The DNN model can be described as a series of functional transformations associated to the model layers (Figure 1).

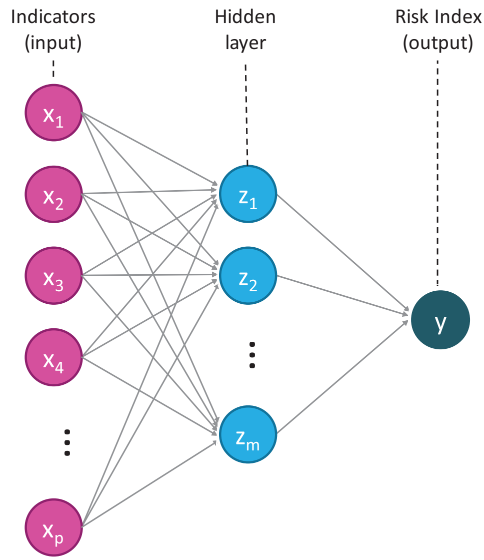

Fig. 1. DNN layers.

The overall length of the chain gives the depth of the model. Specifically, the first network layer performs the following computation of the inputs $X = (x_1, x_2, \dots, x_p)$:

$$a_i = b_i + \sum_{j=1}^{p} x_j w_{i,j} \qquad (2)$$

with $i=1, \dots, m$.

Where a_i, b_i and w_i are respectively defined as activation, bias and model weight.
The activations are transformed by the activation function g within the hidden layer:

$$z_i = g(a_i) \qquad (3)$$

Where z_i is defined as hidden unit. The most used activation function is the sigmoid (Goodfellow et al., 2016). Figure 1 shows only one hidden layer for the sake of simplicity, but there can be several.

The hidden units are combined to give the activations a_o of the output layer:

$$a_o = b_o + \sum_{j=1}^{m} z_j w_{o,j} \qquad (4)$$

Where a_o, b_o and w_o are activation, bias and model weight. Figure 1 shows only one output for the sake of simplicity, but there can be several.

Finally, the activation function h is used to obtain the output y, which is an index for the risk R:

$$y = h(a_o) \sim R \qquad (6)$$

Given a dataset of X and associated y, the model can be trained to minimize the final loss function in a supervised way, in order to predict y based on new inputs X.

3. Application

The described approach was applied to a database of past accidents with the purpose to simulate its application on the national databases managed by the Seveso-competent authorities. The dataset used is the Major hazard incident database (MHIDAS) (AEA technology - Major hazards assessment unit, 2003) launched by the UK Health and Safety Executive in 1986 and developed by AEA Technology until mid 1990's. The events included are based on public domain information sources and their characteristics are registered using keywords.

MHIDAS collects about 8972 hazardous events from 1916 to 1992, recorded by means of the set of items listed in Table 2. Some items use a taxonomy to systematically categorize the event.

Table 2. Set of items used to record hazardous events in MHIDAS (AEA technology - Major hazards assessment unit, 2003). Specific keywords are used to describe some of the items.

Items	**Description**	**Category from taxonomy**
Date	Date of the event	
Location	Location of event	
Substance	Substances involved in the event	X
Event type	Typology of event	X
Origin	Area of the plant and type of equipment from which the event started	X
Section	Plant section in which the event occurred	X
Quantity	Amount (ton) of released substance	
General causes	General causes the led to the event	X
Specific causes	Specific causes the led to the event	X
Evacuated	Number of people evacuated	
	Consequences	
Damage	Economic damage to the property or production loss	
Injured	Number of people injured by the event	
Killed	Number of people killed by the event	

The items listed in the upper part of Table 2 where considered as inputs X to the models, in order to predict the consequences – lower part of Table 2. The details of data pre-processing are explained elsewhere (Solini, 2017). The study focused on the number of people killed and aimed to predict the occurrence of a hazardous event within one of the severity categories listed in Table 3 based on the considered inputs. Only categorical data are used.

Table 3. Severity categories considered by the study.

Severity categories	
0	Event with no fatalities
1-10	Event with a number of fatalities between 1 and 10
10-100	Event with a number of fatalities between 10 and 100

Two datasets were created from the overall MHIDAS database:

1. A training dataset used to train the MLR and the DNN models, with 2/3 of the x_i and associated y values, and
2. A test dataset used to test the models, with about 1/3 of the x_i and associated y values.

A code in Python language was written for training and testing. The classifiers tf.contrib.learn.LinearClassifier and tf.contrib.learn.DNNClassifier from the open-source library TensorFlow (Google LLC, 2018) were used for the models. The DNN model structure (i.e. number of layers and nodes) was inspired by Cheng et al. (2016), based on which the hyper-parameters are defined.

4. Results

The results show whether the events from the test dataset were predicted within the correct severity category from Table 3. The models produce a probability of belonging to a severity category. The probability threshold based on which the decision on whether an event belong to a category is set at 0.5 by default. This affects the following cases:

- true positive (t_p,), as correct prediction of event belonging to a severity category;
- false positive (f_p), as incorrect prediction of event belonging to a severity category;
- true negative (t_n), as correct prediction of event not belonging to a severity category; and
- false negative (f_n), as incorrect prediction of event not belonging to a severity category.

In order to obtain an overall evaluation of the MLR and DNN prediction capabilities, the specific metrics listed in Table 4 were considered. While accuracy, precision and recall are defined based on the default threshold value, the area under the precision/recall curve (PR AUC) is calculated variating the threshold value from 0 to 1.

The results obtained generally show good capability to predict a hazardous event without fatalities in both the methods, as all the metrics reach values at about 0.8 or more.

Table 4. Metrics describing the prediction capabilities of the models.

	Definition
Accuracy	$Acc = \frac{t_p+t_n}{t_p+t_n+f_p+f_n}$
Precision	$Pr = \frac{t_p}{t_p+f_p}$
Recall	$Re = \frac{t_p}{t_p+f_n}$
PR AUC	The area under the Precision/Recall curve.

However, the prediction capabilities sensibly decrease in case of prediction of hazardous events within the severity categories involving fatalities. The only metric increasing for these categories is the accuracy, which almost reaches the unitary value for the category "10 to 100 fatalities". Precision and recall show values under 0.2 and next to 0 respectively for the categories "from 1 to 10 fatalities" and "from 10 to 100 fatalities". PR AUC maintains slightly higher values in both the categories involving fatalities.

5. Discussion

The results show the capability of two machine learning models to manage heterogeneous data from past accident records and extract important information to support safety-related decision making. In fact, the records of hazardous events reported on MHIDAS include both items described by a set of specific keywords and numerical values.

These data were used to build two parallel models predicting the severity of potential new hazardous events. The metrics obtained from testing the two models show good capabilities in predicting hazardous events without fatalities. However, it must be pointed out the presence of class imbalance as the events without fatalities represent the vast majority within the dataset considered. This could be the reason of such a variation in performance when we look at the remaining severity categories.

Hazardous events that cause from 1 to 10 fatalities were predicted with an accuracy that is over 0.8, while the category "from 10 to 100" show even higher accuracies. However, it must be considered that the accuracy metric represents a partial evaluation of the model. In fact, if the model is employed for the prediction of rare events (such as the ones with fatalities), predicting always their "non-occurrence" would lead to high accuracy anyway as the metric presents the term of "true negatives" at the numerator. This is demonstrated by the other metrics, in particular precision and recall, which are equal to 0 in the last category reflecting the absence of predicted "true positives". For this reason, the evaluation of the model capabilities can be carried out only through the whole set of metrics.

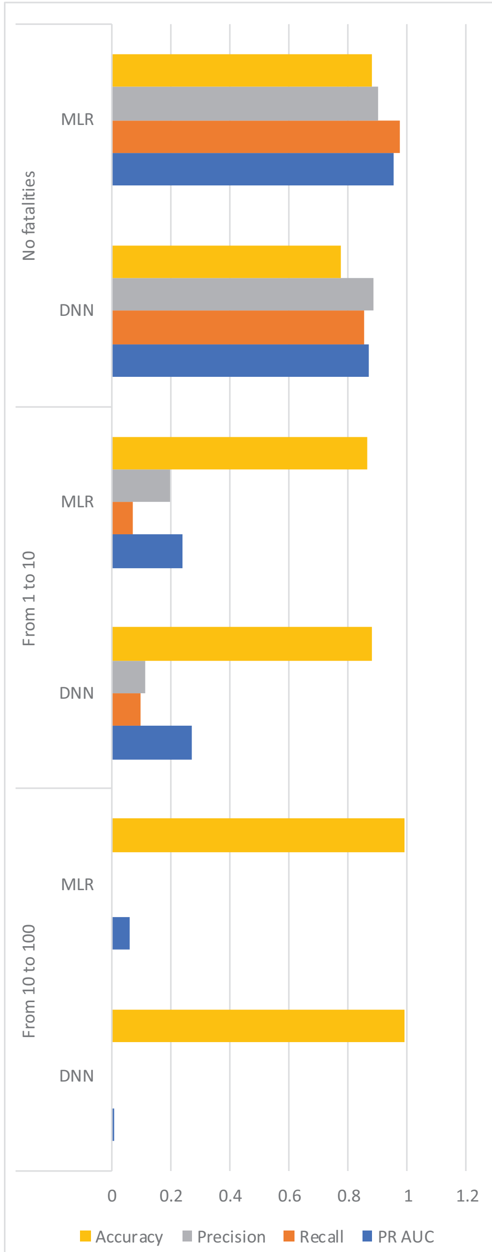

Fig. 2. Test results for the multiple linear regression (MLR) and deep neural network (DNN) models.

PR AUC gives a more complete overview of the model as it may indicate also the potential for improvement. In fact, it indicates the possibility to adjust the decision probability threshold to improve precision or recall based on the purpose of the analysis. For instance, tuning the threshold to optimize the model precision would be effective when predicting a rather frequent event with relatively low criticality, such as a hazardous event without fatalities. This would allow obtaining a model that tends to avoid false positives further increasing the number of alarms with some false ones. On the other hand, when predicting a rare event with relatively high criticality (such as a high-impact low-probability accidents (Pasman, 2015)), false alarms may be tolerated in exchange of a better prediction of the actual hazardous event, thus an improvement in true positive. For this reason, the threshold may be adjusted to optimize the recall.

A comparison between the two models shows a relatively better performance in the predictions obtained by MLR. However, if we compare the precision values for the severity category without fatalities, DNN turns to be more appropriate. Analogously, comparing the recall values for the severity category with fatalities between 10 and 100 shows a relatively better performance by DNN.

This demonstrates that there are some important differences among the specific techniques. Linear models such as MLR are widely used for prediction purposes. Interactions of the event features can be easily memorized through the provided datasets. However, a relatively simple model may not be able to capture the essential pattern in the data (Christian and Griffiths, 2016). Generalization of lessons learned for prediction under unknown circumstances requires a higher level of complexity, which linear functions may fail to provide (Goodfellow et al., 2016). Deep neural networks may be an option for such task (Christian and Griffiths, 2016).

Major accidents are (fortunately) rare events in industry, even considering evidence of fat-tailed distributions (Taleb, 2007). For this reason, appropriate models should be used to deal with such unexpected events. To this end, linear regression techniques are well-known for their limitation to handle rare events data (King and Zeng, 2001). Relatively simple models tend to forecast the basic trend and may potentially miss several exact points (Christian and Griffiths, 2016). Sophisticated models such as DNN are theoretically better suited to consider rare events, due to their sensitivity to input data and capability to generalize (Cheng et al., 2016). However, a limitation of DNN is that the model may have such a sensitivity to input data that the solutions it produces are highly variable (Christian and Griffiths, 2016). There can be errors in how the data were collected or reported on MHIDAS. For this reason, cross-validating with a test dataset is essential. Moreover, DNN results can be altered by its random initialization of parameters before every training session. This has the potential to affect the whole model development and, in turn, lead to slight alterations of prediction capabilities. Such differences may be amplified in case of relatively small datasets and few iterations to minimize the final loss function during training. Another limitation of the DNN model used in this case study may be related to its setting based on Cheng et al.'s (2016) work. In fact, the DNN model used may still need appropriate optimization for the case study.

An important aspect to consider is that the DNN model is not tied to a rigid structure to aggregate information from indicators (Landucci and Paltrinieri, 2016), but it has the potential to reshape its own structure based on new batches of data. Such an approach has some similarity with other methodologies in literature (Paltrinieri et al., 2016, 2015b), who developed a technique to update logic trees describing accident scenarios dynamically, in order to account for new evidence and prevent emergence of atypical events.

An option for improving the DNN model is the application of progressive learning techniques, which may be independent of the number of indicator categories and to learn new indicators once relevant information emerges, while retaining the knowledge of previous ones (Venkatesan and Er, 2016).

6. Conclusions

Through this work, we have suggested a data analytics approach for severity prediction of future hazardous events. The approach was based on the use of two well-known machine learning techniques: MLR and DNN. These models were developed and tested on the Major Hazardous Incident Data Service (MHIDAS) database. Part of the available data were used to build the actual models while the remaining data were used to test the models. This allowed also identifying and discussing the inherent limitations of the techniques.

For instance, DNN high model sensitivity does not tolerate inaccurate inputs. For this reason, selection and customization of a prediction model for an intended purpose should be carefully carried out using appropriate metrics, tolerance, and criteria. If these precautions are considered, the odds to deliver appropriate support for safety-related decision-making will be boosted. In this way, it will be possible to extract important information from heterogeneous data and effectively support safety-related decision making.

Acknowledgement

We thank Riccardo Solini for his work on this topic, which ultimately inspired this contribution.

References

AEA technology - Major hazards assessment unit, 2003. MHIDAS – Major Hazard Incident Data Service. UK.

Andrews, D.F., 1974. A Robust Method for Multiple Linear Regression. Technometrics 16, 523–531. https://doi.org/10.1080/00401706.1974.10489233

Cheng, H.-T., Koc, L., Harmsen, J., Shaked, T., Chandra, T., Aradhye, H., Anderson, G., Corrado, G., Chai, W., Ispir, M., 2016. Wide & deep learning for recommender systems, in: Proceedings of the 1st Workshop on Deep Learning for Recommender Systems. ACM, pp. 7–10.

Christian, B., Griffiths, T., 2016. Algorithms to live by: the computer science of human decisions. Macmillan.

European Parliament and Council, 2012. Directive 2012/18/EU of 4 July 2012 on the control of major-accident hazards involving dangerous substances, amending and subsequently repealing Council Directive 96/82/EC - Seveso III. Off. J. Eur. Union 1–37.

Goodfellow, I.J., Bengio, Y., Courville, A., 2016. Deep learning, The MIT Press. Citeseer, Cambridge, Massachusetts, US.

Google LLC, 2018. Tensorflow [WWW Document]. Tensorflow TM. URL www.tensorflow.org

Hastie, T., Tibshirani, R., Friedman, J., 2009. Unsupervised learning, in: The Elements of Statistical Learning. Springer, pp. 485–585.

King, G., Zeng, L., 2001. Logistic regression in rare events data. Polit. Anal. 9, 137–163.

Landucci, G., Paltrinieri, N., 2016. Dynamic evaluation of risk: From safety indicators to proactive techniques. Chem. Eng. Trans. 53. https://doi.org/10.3303/CET1653029

Øien, K., 2001. A framework for the establishment of organizational risk indicators. Reliab. Eng. Syst. Saf. 74, 147–167. https://doi.org/10.1016/S0951-8320(01)00068-0

Paltrinieri, N., Comfort, L., Reniers, G., 2019. Learning about risk: Machine learning for risk assessment. Saf. Sci. 118, 475–486. https://doi.org/https://doi.org/10.1016/j.ssci.2019.06.001

Paltrinieri, N., Cozzani, V., Wardman, M., Dechy, N., Salzano, E., 2010. Atypical major hazard scenarios and their inclusion in risk analysis and safety assessments, in: Reliability, Risk and Safety: Back to the Future. pp. 588–595.

Paltrinieri, N., Khan, F., 2016. New Definitions of Old Issues and Need for Continuous Improvement, in: Dynamic Risk Analysis in the Chemical and Petroleum Industry: Evolution and Interaction with Parallel Disciplines in the Perspective of Industrial Application. https://doi.org/10.1016/B978-0-12-803765-2.00002-0

Paltrinieri, N., Khan, F., Cozzani, V., 2015a. Coupling of advanced techniques for dynamic risk management. J. Risk Res. 18, 910–930. https://doi.org/10.1080/13669877.2014.919515

Paltrinieri, N., Øien, K., Tugnoli, A., Cozzani, V., 2013. Atypical accident scenarios: From identification to prevention of underlying causes, Chemical Engineering Transactions. https://doi.org/10.3303/CET1331091

Paltrinieri, N., Reniers, G., 2017. Dynamic risk analysis for Seveso sites. J. Loss Prev. Process Ind. 49. https://doi.org/10.1016/j.jlp.2017.03.023

Paltrinieri, N., Tugnoli, A., Bonvicini, S., Cozzani, V., 2011. Atypical scenarios identification by the DyPASI procedure: Application to LNG. Chem. Eng. Trans. 24, 1171–1176. https://doi.org/10.3303/CET1124196

Paltrinieri, N., Tugnoli, A., Cozzani, V., 2016. Dynamic Hazard Identification: Tutorial and Examples, in: Dynamic Risk Analysis in the Chemical and Petroleum Industry: Evolution and Interaction with Parallel Disciplines in the Perspective of Industrial Application. https://doi.org/10.1016/B978-0-12-803765-2.00004-4

Paltrinieri, N., Tugnoli, A., Cozzani, V., 2015b. Hazard identification for innovative LNG regasification technologies. Reliab. Eng. Syst. Saf. 137. https://doi.org/10.1016/j.ress.2014.12.006

Pasman, H.J., 2015. Risk Analysis and Control for Industrial Processes - Gas, Oil and Chemicals: A System Perspective for Assessing and Avoiding Low-Probability, High-Consequence Events. Elsevier Science.

Paté-Cornell, E., 2012. On “Black Swans” and “Perfect Storms”: Risk Analysis and Management When Statistics Are Not Enough. Risk Anal. 32, 1823–1833. https://doi.org/10.1111/j.1539-6924.2011.01787.x

Samuel, A.L., 1959. Some Studies in Machine Learning Using the Game of Checkers. IBM J. Res. Dev. 3, 210–229. https://doi.org/10.1147/rd.33.0210

Solini, R., 2017. Data analytics for chemical process risk assessment: learning lessons from past events towards accident prediction. Bologna, Italy.

Sornette, D., 2009. Dragon-Kings, Black Swans and the Prediction of Crises. ETH Zurich, Chair of Systems Design.

Svozil, D., Kvasnicka, V., Pospichal, J., 1997. Introduction to multi-layer feed-forward neural networks. Chemom. Intell. Lab. Syst. 39, 43–62. https://doi.org/10.1016/S0169-7439(97)00061-0

Taleb, N., 2007. The black swan : the impact of the highly improbable. Random House, New York.

Venkatesan, R., Er, M.J., 2016. A novel progressive learning technique for multi-class classification. Neurocomputing 207, 310–321.