



ELSEVIER

Contents lists available at ScienceDirect

Safety Science

journal homepage: www.elsevier.com/locate/safety

Learning about risk: Machine learning for risk assessment

Nicola Paltrinieri^{a,*}, Louise Comfort^b, Genserik Reniers^{c,d,e}^a Department of Mechanical and Industrial Engineering, Norwegian University of Science and Technology – NTNU, Norway^b Center for Disaster Management, Graduate School of Public and International Affairs, University of Pittsburgh, USA^c Faculty of Applied Economics, University of Antwerp Operations Research Group ANT/OR, Antwerp, Belgium^d Center for Corporate Sustainability (CEDON), HUB, KULeuven, Brussels, Belgium^e Safety Science Group, TU Delft, Delft, the Netherlands

ARTICLE INFO

Keywords:

Risk assessment
Dynamic risk analysis
Machine learning
Deep learning

ABSTRACT

Risk assessment has a primary role in safety-critical industries. However, it faces a series of overall challenges, partially related to technology advancements and increasing needs. There is currently a call for continuous risk assessment, improvement in learning past lessons and definition of techniques to process relevant data, which are to be coupled with adequate capability to deal with unexpected events and provide the right support to enable risk management. Through this work, we suggest a risk assessment approach based on machine learning. In particular, a deep neural network (DNN) model is developed and tested for a drive-off scenario involving an Oil & Gas drilling rig. Results show reasonable accuracy for DNN predictions and general suitability to (partially) overcome risk assessment challenges. Nevertheless, intrinsic model limitations should be taken into account and appropriate model selection and customization should be carefully carried out to deliver appropriate support for safety-related decision-making.

1. Introduction

Shifts in our understanding of risk are continuously imposed by the emergence of new knowledge, reshaping the boundaries of our definitions. This is particularly important in safety-critical sectors, such as the petroleum and chemical industry, often striving for enhanced system performance, but where unwanted and related losses of hazardous substances can endanger a large number of people at once (Paltrinieri and Khan, 2016a,b).

One of the most renowned definitions of risk was given by Kaplan and Garrick (1981). It states that risk (R) can be expressed by what can go wrong (scenario s), what likelihood it will have (probability p), and how severe consequences will be (consequence c):

$$R = f(s, p, c) \quad (1)$$

Since Kaplan and Garrick, numerous attempts have been made by analysts and scholars to capture the notion of risk in a more meaningful way. Occurrence of events that “went seriously wrong”, such as major industrial accidents, are unfortunate reminders of the details that cannot be framed by formula 1 (Paltrinieri et al., 2012a). Aven (2012) provides a thorough review of the nuances of (and the lack of a common agreement on) the risk concept.

Recognizing the high costs and consequences of large-scale

industrial accidents, organizational theorists searched for methods to ensure high reliability in organizational performance in risk environments (LaPorte and Consolini, 1991; Rochlin et al., 1987; Roe and Schulman, 2008; Weick and Roberts, 1993; Weick and Sutcliffe, 2001). Some facilities (nuclear power plants, aircraft carriers, high-speed trains, hazardous materials storage sites) provide benefit to society, but are inherently risky. For these large-scale sociotechnical systems, failure has an increased criticality due its potential cost in lives, equipment, and destruction to the community. The challenge is how to ensure human control over technical operations that are potentially dangerous.

The prescribed means to ensure highly reliable performance in risk conditions include defining a clear set of rules of operation for uncertain contexts, designing advanced training for managing the equipment and tasks involved, and fostering practice of heedful interaction among actors and components of the operating environment for anomalies that may indicate potential threats (LaPorte and Consolini, 1991; Weick and Roberts, 1993). The capacity to produce highly reliable performance depends upon deep knowledge of the operating environment and its limitations, intensive communication among participants, and acceptance of a culture in which each member of the organization accepts responsibility for correcting observed errors in any part of the organization’s performance (Hutchins, 1995).

* Corresponding author.

E-mail address: nicola.paltrinieri@ntnu.no (N. Paltrinieri).<https://doi.org/10.1016/j.ssci.2019.06.001>

Received 12 July 2018; Received in revised form 31 May 2019; Accepted 1 June 2019

0925-7535/© 2019 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Yet, as the technologies of sensors and timing advanced in managing high risk operations, the focus on organizational control shifted to highly sensitive programs of computational management of machines that combined multiple measures of performance to provide more consistently reliable management of technical operations in changing risk environments (Nobre, 2009). The risk remained, but the management practice and technologies changed.

Villa et al. (2016a, 2016b) demonstrate how different risk definitions may affect the approach taken for its assessment and management. Villa et al. (2016a, 2016b) also remind us that, while quantitative risk assessment (QRA) is required by law in several industrial sectors, it is performed mainly during the design phase. For this reason, it only describes a static risk picture of the system (Pasman and Reniers, 2014). The issue of realistically evaluating a given scenario s is also addressed by Apostolakis (2004) and Creedy (2011). They question the probabilities and frequencies used in quantitative risk analysis, affirming that they are retrieved from outdated databases and they may not fit the studied system. They also affirm that probability calculation is heavily affected by scarcity of data. Landucci et al. (2016a, 2016b) demonstrate how the impact of an unwanted event is influenced by a series of dynamic variables, which are not always considered for its prediction. Moreover, if we want to assess the overall risk covering all the possible scenarios $s_i, i = 1, \dots, N$, how do we know that we are not missing anything and $N = N_{max}$? We cannot be sure that we will be free from “atypical” scenarios, as theorized by Paltrinieri et al. (2013, 2012a); that is scenarios that are not captured by standard hazard identification techniques because they deviate from normal expectations of unwanted events or worst-case scenarios.

This study proposes a solution to the risk assessment main challenges based on the application of machine learning techniques. While the following section introduces the additional risk dimension of knowledge and summarizes the state of the art of the industrial risk assessment main challenges, section 3 describes indicator-based approaches and a representative case study from the offshore Oil & Gas industry. Machine learning and Deep Neural Networks (DNN) are suggested as a possible solution and applied to the case study in section 4. Section 5 illustrates application results, section 6 discusses benefits and limitations of machine learning for risk assessment, and section 7 provides some conclusions.

2. Risk knowledge

Aven and Krohn (2014) suggest including a new dimension in the definition of risk (R): knowledge (k):

$$R = f(s, p, c, k) \tag{2}$$

Fig. 1a shows how a two-dimensional risk matrix depicts formula 1. A traffic-light colour code represents acceptable (green), unacceptable (red) or intermediate (yellow and orange) risk. The application of the additional knowledge dimension (formula 2) would bend the matrix as depicted in Fig. 1b. Expressing the level of knowledge used for risk assessment is an intrinsic feature of the calculated value of risk. This implies the definition of a condition of unacceptable knowledge, which may be represented by the space under the matrix in Fig. 1b. We can tolerate having relatively little knowledge of scenarios with both low probability and low consequence. For this reason, the matrix is bent towards its minimum values in this area. The matrix reaches its peak where red is more intense and probability and consequence have their highest values. This represents the need for thorough knowledge of scenarios falling in this area.

Formula 2 gives important insight on how we should treat risk assessment results and supports the continuous improvement of the analysis – we become aware of how uncertainty is always a companion and that we should cope with it (De Marchi and Ravetz, 1999). For this reason, we adopt this formula as the basis for this study among numerous definitions of risk (Aven, 2012). However, another question emerges: how can we consider knowledge in quantitative risk assessment? In addition, even if we can assess risk with all the knowledge available, we would provide a risk picture that is “frozen” in time, while the system is changing around it. The conditions considered on day zero may not be valid anymore on day one. For this reason, we also need to address how to consider system evolutions. Calibration and correction based on new evidence would possibly allow risk analysis to consider evolving conditions and improve system knowledge. Such a dynamic approach to risk management is theorized and reviewed by a number of studies (Khan et al., 2016; Paltrinieri et al., 2014; Paltrinieri and Khan, 2016a,b; Villa et al., 2016a).

Underlying the dynamic approach to risk management is the concept of “initial conditions” that set the trajectory for evolving system performance (Kaufmann, 1993; Prigogine and Stengers, 1984). Initial conditions represent the existing state of an organization at risk, prior to a specific hazardous event. It includes the basic resources available for learning and action, as well as the current operating context of the organization. These conditions shape the possible courses of action for coordinated response to an actual event (Comfort, 2019, 1999). Given the distinctive set of initial conditions, an organization engages in an evolving learning process that reflects its practical response to risk, its interaction with other organizations and conditions, and produces the next (temporary) state of operations. The set of interactive responses by organizations with the environment, repeated over time, constitutes a dynamic response system as it adapts to risk.




Fig. 1. (a) Two-dimensional risk matrix according to formula (1); (b) three-dimensional risk matrix according to formula (2).

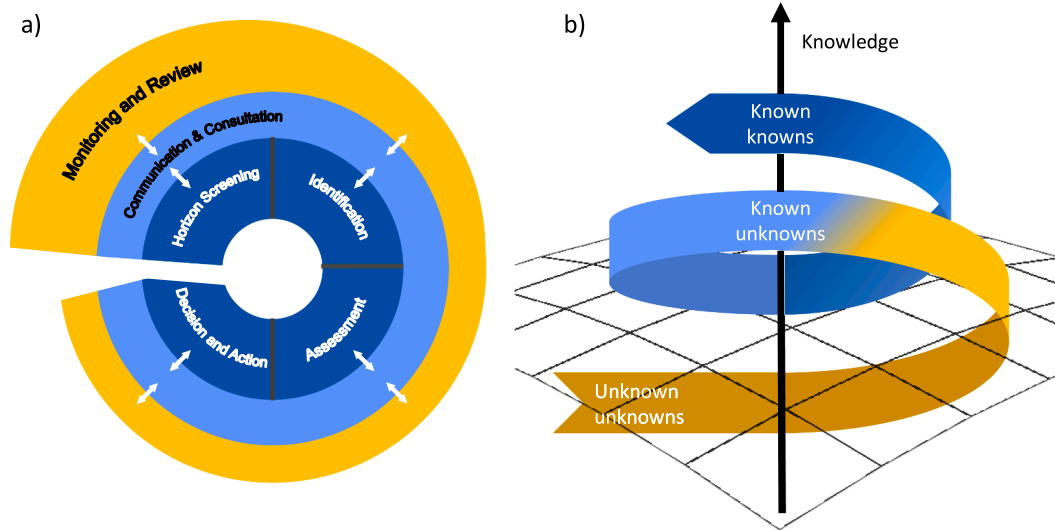


Fig. 2. (a) Dynamic Risk Management Framework (DRMF - clockwise), adapted from Paltrinieri et al. (2014); (b) DRMF revolving around the knowledge dimension, adapted from Villa et al. (2016a).

Fig. 2a represents the Dynamic Risk Management Framework (DRMF) defined by Paltrinieri et al. (2014). DRMF focuses on continuous systematization of information on new risk evidence. Its shape is open to the outside to avoid vicious circles and self-sustained processes. It opens the process to new information, early warnings and unwanted events by means of continuous monitoring. Such information is an input (through communication or consultation) to each of the four steps of risk management. There is no end to the process, but iteration, in order to keep track of changes and elaborate them for improved management. Such iteration is in accordance with the revised definition of risk in formula 2, as shown by the three-dimensional representation of DRMF (Fig. 2b) revolving around the dimension of knowledge to escape the aforementioned space of unacceptability.

Epistemic limitations and continuous modifications of the world around us lead to an obvious conclusion: there will be always something that we cannot capture while assessing risk. Within the space of unacceptable knowledge we may encounter Unknown unknown events (as defined in Table 1), or Black Swans. Taleb (2007) defines such events as those that can be explained only after the fact and cannot be anticipated. Our best chance to lower risk is being aware that there are scenarios that we do not know (in part or at all – Known unknowns in Table 1) and implement DRMF. This represents a way out from unacceptable knowledge towards Known knows (Table 1). Nevertheless, knowledge may be disregarded or simply forgotten, covering the spiral in Fig. 2b backward and incurring in Unknown knows (Table 1). This underlines that fact that the main challenge is effectively capitalizing the accumulated knowledge and avoiding its oblivion.

Nowadays, emerging cyber-physical systems within industry present a significant opportunity to implement DRMF. Such systems embed internet of things solutions and wireless sensor networks, allowing for collection of data records in all phases of product lifecycle (Lasi et al., 2014; Wang et al., 2016). Lasi et al. (2014) state that the increasing digitalization in industry is resulting in the registration of an increasing amount of actor- and sensor-data which can support functions of

dynamic risk analysis, as opposed to traditional risk analysis incapable of reflecting evolving real-world risk (Paltrinieri and Khan, 2016a,b; Yang et al., 2017). However, increasing complexity creates uncertainty about technological capabilities and adequate strategies to apply them (Schumacher et al., 2016). For this reason, the transformation of risk models should result in handy software tools to enable DRMF application in practice.

2.1. State of the art and overall challenges

A number of approaches address the need of continuous update of risk assessment and may be grouped in two macro groups: empirical and theoretical. First-group approaches are generally developed by observing a large amount of relevant data. Whereas, sparse data would lead to relying on theory-based approaches – given some inevitable assumptions. Fig. 3 depicts an overall simplification of the state of the art of risk assessment and the ideal risk assessment approach on a models/data graph.

Representative examples of approaches from industrial applications rely on simplified (empirical) models and a big amount of data, as reported in the following.

- Popular software for bowtie analysis allow for real-time monitoring of safety measures performance (CGE Risk Management Solutions B.V., 2016). End-users are from Oil & Gas, aviation, mining, maritime industries and healthcare may use such software to support risk management.
- Attempts to better monitor safety measures were carried out by the Norwegian Oil & Gas industry (Etterlid and Etterlid, 2013; Hansen, 2015; Statoil, 2013). However, they only rely on monitoring and several of them have been suspended during the 2014–2015 oil crisis.
- Preliminary methodologies developed in collaboration with industry can also be found in literature (Risk Barometer (Hauge et al., 2015)

Table 1
Definitions of Known/Unknown events (Paltrinieri et al., 2012a).

Unknown unknowns	Known unknowns	Unknown knows	Known knows
Events we are not aware that we do not know, whose risk cannot be managed	Events we are aware that we do not know, for which we employ both prevention and learning capabilities	Events we are not aware that we already know, or used to know, with certain confidence	Events we are aware that we know, whose risk we can manage with a certain level of confidence

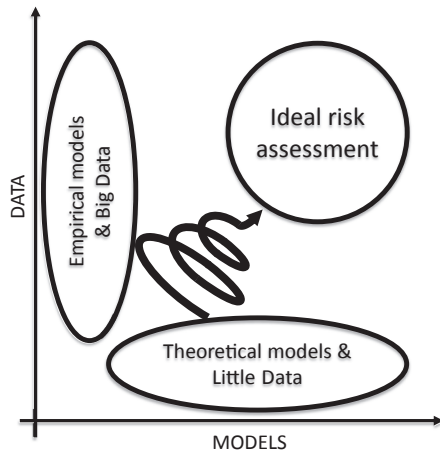


Fig. 3. Graphical representation of the state of the art of risk assessment and the ideal risk assessment approach on a models/data graph.

and Modelling Instantaneous Risk for Major Accident Prevention – MIRMAP (Yang and Haugen, 2015)).

Risk assessment models with relatively strong theoretical bases may be found in literature and mainly aim at dynamic risk assessment. Nevertheless, lack of data from real cases has led to large sets of assumptions and simulations for their development. Representative examples in the following.

- Khakzad has extensively worked on the application of Bayesian networks to dynamic risk assessment problems in the chemical process industry (Khakzad, 2015; Khakzad et al., 2014, 2013a).
- Several contributions to dynamic risk assessment by means of the Monte Carlo method can be found in literature (Noh et al., 2014; Targoutzidis, 2012). Such contributions are either applied (Noh et al., 2014) or addressing new findings on a purely methodological level (Durga Rao et al., 2009; Targoutzidis, 2012).
- The Petri nets method is also used to improve risk assessment and capture dynamic sequences (Nivolianitou et al., 2004; Nývlt et al., 2015; Nývlt and Rausand, 2012; Zhou et al., 2017; Zhou and Reniers, 2017, 2016a, 2016b).

Improving risk assessment would mean to iteratively learn from this experience and provide an ideal approach that relies on both Big Data and theoretical models. A first effort to move from the ellipses in the Models/Data graph (Fig. 3), is provided by Paltrinieri and Khan (2016a,b). However, we can summarize five main methodological challenges to be addressed in such a journey towards “ideal” risk assessment:

1. **Dynamicity:** how do we continuously update and improve risk assessment? It would allow refining the considered set of possible accident scenarios with accurate likelihood and impact. This would keep track of change and evolution of the industrial system.
2. **Cognition:** how do we learn from relevant lessons to improve risk assessment? Unwanted events and experts can provide valuable insight. Capitalizing such knowledge in a systematic way would prevent accident repetition.
3. **Data processing:** how do we process big data from the industrial system? Risk modelling should develop appropriate manipulation of the large datasets collected today in industry, because they describe system state and would produce meaningful risk information.
4. **Emergence:** how do we prepare for what we do not know? This challenge refers to the need of addressing emerging (not known before) risks. This is fundamental in relation to new technologies on which there is relative lack of risk experience.
5. **Usability:** how do we provide a real support and allow for implementation of lessons learned in industry? This last challenge reflects practical industry needs. It refers to the need of comprehensive support to decision-making.

3. Small things

Paté-Cornell (2012) and Haugen and Vinnem (2015) warn against the misuse of the Black Swan concept. This should not be a reason for ignoring potential scenarios or waiting until a disaster happens, to take safety measures and issue regulations against a predictable situation. On the contrary, it should represent an incentive to continuously learn and improve (as suggested by Fig. 2b). What can we do against what we do not know? Sornette (2009) provides an answer to such concern by applying a geophysical model (Musgrave, 2013) on the prediction of earthquakes. He saw that some degrees of organization and coordination could serve to amplify small fractures, always present and forming in the tectonic plates. Organization and coordination may turn small causes into large effects, i.e. large earthquakes characterized by low probability. Paltrinieri and Khan (2016a,b) are in line with this, claiming that extreme accidents may be described as a particular combination of single events, some of which may be considered as “Small Things” – e.g. apparently meaningless technical malfunction or human distraction. Acting on Small Things would allow breaking the chain of events leading to an accident and lowering its probability.

A number of approaches are used to describe accident sequences and understand how to stop them. Some of the most known and used in industry are logic trees such as fault tree, event tree and bow-tie diagram (Center for Chemical Process Safety, 2000). An example from the offshore Oil & Gas industry is shown in Fig. 4 and further described in Section 3.1. Logic trees are used to evaluate risk on a probabilistic basis. The concept of “safety barriers” is used to model and include prevention and/or mitigation measures. The Norwegian oil & gas sector (Petroleum Safety Authority, 2013) commonly uses a specific hierarchical structure




Fig. 4. Event tree describing a drive-off scenario of an Oil & Gas drilling rig. EDS stands for Emergency Disconnect System. Adapted from (Paltrinieri et al., 2016b).

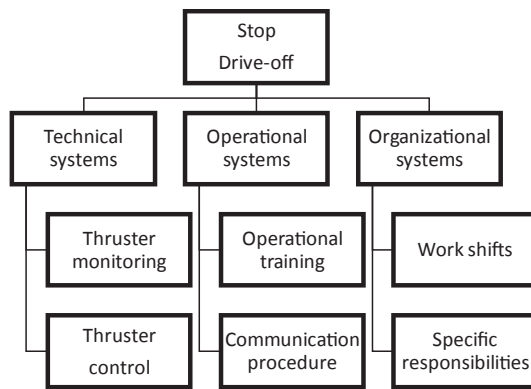


Fig. 5. Hierarchical structure of the safety barrier “stop drive-off”.

to model safety barriers (Fig. 5), defining them as “systems of technical, operational and organisational elements, which are intended individually or collectively to reduce the possibility for a specific error, hazard or accident to occur, or which limit its harm/disadvantages” (Petroleum Safety Authority, 2013).

In the last decade, increasing attention has been dedicated to monitoring and evaluation of safety barrier performance through indicators, as a way to assess and control risk. Indicators may report a series of factors: physical conditions of a plant (equipment pressure and temperature), number of failures of an equipment piece, maintenance backlog, number of emergency preparedness exercises run, amount of overtime worked, etc. They overlap with the concept of Small Things. A number of indicator typologies are theorized and used in literature (Øien et al., 2011). Øien et al. (2011) affirm that we can refer to risk indicators if: 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. That being said, the latter feature is quickly becoming outdated due to the extensive collection carried out in industry and the attempts to process large numbers of them (Paltrinieri and Reniers, 2017).

Øien et al. (2011), Paltrinieri et al. (2016a, 2012b), and Landucci et al. (2016) have produced several reviews on risk and barrier indicators. They show that definition and collection of risk indicators have become consolidated practices in “high-risk” sectors, such as the petroleum and chemical industries. For instance, the Norwegian Petroleum Safety Authority (PSA) requires indicators describing the technical performance of safety barriers within the Norwegian Oil & Gas industry since 1999 (PSA, 2016); while, the European directive “Seveso III” (European Parliament and Council, 2012) on the control of major-accident hazards involving dangerous substances suggests their use for sites handling hazardous substances (European Parliament and Council, 1982). Such trend towards definition and collection of higher numbers of indicators (Paltrinieri and Reniers, 2017) demonstrates the mentioned challenge on big data process for risk level assessment.

3.1. Representative case study from the offshore Oil & Gas industry

In order to avoid potential damage during drilling operations for a new offshore Oil & Gas well, a semi-submersible drilling unit should maintain the position above the wellhead. This is particularly critical if the platform is located in shallow waters, where small changes of position lead to higher riser (pipe connecting the platform to the subsea drilling system) angles. Exceeding physical inclination limits may result in damages to wellhead, Blowout Preventer (BOP – sealing the well) or Lower Marine Riser Package (LMRP – connecting riser and BOP) (Chen et al., 2008). Platform positioning is maintained in an autonomous way (without mooring system) through the action of a set of thrusters controlled by the Dynamic Positioning (DP) system. Input for the DP system is provided by the position reference system (Differential Global

Positioning System – DGPS and Hydroacoustic Position Reference – HPR), environmental sensors, gyrocompass, radar and inclinometer (Chen et al., 2008). A Dynamic Positioning Operator (DPO) located in the Marine Control Room (MCR) is responsible for constant monitoring of DP panels and screens and carrying out emergency procedures if needed (Giddings, 2013).

Platform position may be lost due to a series of reasons. In this case study, it is assumed that the platform thrusters exercise propulsion towards a wrong direction, leading to a scenario of “drive-off. If the rig moves to an offset position, specific alarms turn on and suggest the DPO to stop the drive-off scenario by deactivating the thrusters and initiating the manual Emergency Disconnect Sequence (EDS) for the disconnection of the riser from the BOP. If the manual EDS ultimately fails, the automatic EDS activates at the ultimate position limit allowing for safe disconnection (Chen et al., 2008). Fig. 4 depicts the event tree describing such a drive-off scenario (Paltrinieri et al., 2016b).

Matteini (2015) studies in detail occurrence and development of drive-off scenarios. She modelled each safety barrier reported in the event tree of Fig. 4 with the hierarchical structures of their technical, operational and organizational systems. Relevant indicators are also defined to assess the performance of systems, and, in turn, barriers. Matteini defines 50 indicator categories in total, whose values are to be collected and translated on a mutually comparable criticality scale, ranging between 1 and 6. Indicator trends are simulated for a period of 30 years. They are inspired to the typical bathtub curve for technical elements (Wang et al., 2002) and relevant expert judgment for the remaining elements. Fig. 6 reports examples of indicator trends that are associated to the barrier systems in Fig. 5. As shown by Bucelli et al. (2017), indicator values may be aggregated based on relative weights and hierarchical barrier models, in order to enable dynamic update of barrier failure probabilities. This can be used to update, in turn, occurrence frequencies of potential outcomes by reiteratively performing the event tree analysis, as shown by the frequency of wellhead damage in Fig. 6. Outcome frequencies are an expression of the scenario probability p mentioned in (1) and, in turn, of the risk R , if we assume that the other factors are constant. Matteini (2015) points out a certain complexity within the model, which may be due to a tangled structure and an unclear approach to assign relative weights to single model elements.

4. Machine learning

Diekmann (1992) stated: “new analysis tools are emerging, which have the potential to allow complex risk analyses to be performed simply. These new tools, which are underpinned by decision analysis and, lately, expert-systems technology, may lead to powerful, yet simple, approaches to the representation of risky problems.” He also suggested a possible interdisciplinary direction for the evolution of risk analysis by stating: “future approaches to risk analysis will certainly rely more on the advances being made in artificial intelligence and the cognitive sciences. New computer tools and knowledge-representation schemes will unquestionably lead to new techniques, insights and opportunities for risk analysis.”

However, industrial risk analysis has unevenly progressed since this statement, not respecting Diekmann’s prediction and leaving a series of methodological gaps, as shown in Section 2.1. At the same time, the use of artificial cognition has possibly become more attractive, given the progressive refinement of its models and the exponential increase in available computing power (Goodfellow et al., 2016).

This study suggests solutions from a branch of machine learning denominated “deep learning” and shows how these can address some of the risk analysis gaps. Machine learning refers to techniques aiming to program computers to learn from experience (Samuel, 1959). Deep learning 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

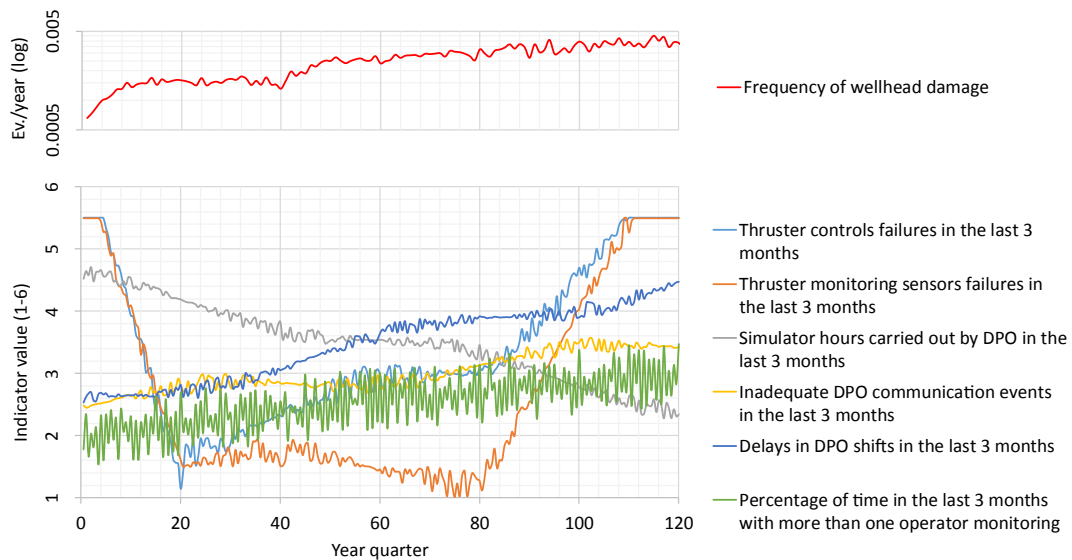


Fig. 6. Simulated frequency trend of wellhead damage and examples of simulated indicator trends associated to the barrier systems in Fig. 5, according to Matteini’s (2015) study.

layers to learn representations of data with multiple levels of abstraction.

A computer may be trained to assess risk for safety-critical industries such as Oil & Gas through deep learning techniques (Fig. 7). This would allow processing a large amount of information in the form of indicators from normal operations and past unwanted events (from mishaps to major accidents), which would be used for training. Due to the subjectivity of risk definition (as discussed in section 2) risk level cannot be assigned to each event with certainty and expert supervision is needed. Deep learning allows for this supervised learning (Goodfellow et al., 2016). Once the model has learned risk categorization, it uses its knowledge to evaluate real-time risk from the state of

the monitored system, e.g. an offshore Oil & Gas platform.

4.1. 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). The model was chosen due to its similarity with the hierarchical structure used to aggregate indicator information (Fig. 5). A linear model, such as a linear regression, would be restricted to linear functions, while a DNN model describes the target as a non-linear function of the input features (Goodfellow et al., 2016). The DNN model can be described as a series of functional transformations




Fig. 7. Representation of deep learning approach for risk assessment.

associated to the model layers (Fig. 7). The overall length of the chain gives the depth of the model. The name “deep learning” derives from this (Goodfellow et al., 2016). Specifically, the first network layer performs the following computation of the inputs x_b, \dots, x_p , which, in this case, are performance indicators:

$$a_i = b_i + \sum_{j=1}^p x_j w_{i,j} \tag{3}$$

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:

$$z_i = g(a_i) \tag{4}$$

where z_i is defined as hidden unit. The most used activation function is the sigmoid (Goodfellow et al., 2016). Fig. 5 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:

$$a_o = b_o + \sum_{j=1}^m z_j w_{o,j} \tag{5}$$

where a_o, b_o and w_o are activation, bias and model weight. Fig. 7 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, in this case, is an index for risk R :

$$y = h(a_o) R \tag{6}$$

Given a dataset of x_i and associated y , the model can be trained to minimize the final loss function in supervised way (Goodfellow et al., 2016), in order to predict y based on new inputs x_i .

4.2. Model application

Matteini (2015) has simulated the trend of 50 different indicator categories over 30 years (Fig. 6) to assess the performance of the safety barriers involved in a drive-off scenario (Fig. 4). Indicator readings are assumed every 6 weeks for a total of 240 values per indicator category. As already mentioned, aggregation of these indicators through relatively complex barrier hierarchical structures and event tree analysis allowed assessing the wellhead damage frequency over time (Fig. 6). Trend definition is particularly important in terms of decision-making support, because it allows the operator to understand whether the system is improving or worsening in terms of risk. For this reason, the study focuses on the prediction of risk increase given the indicator trends.

Since the simulated wellhead damage frequency $Freq$ is an expression of the scenario probability p , and, in turn, the risk R , for constant scenario s and consequence c , we can state that:

$$\frac{dFreq}{dt} \approx \frac{dR}{dt} \tag{7}$$

For this reason, $Freq$ was transformed into its derivative with respect to time t , and labels indicating its increase or decrease were added within the database (Table 2).

Table 2
Definition of the output used as risk index to predict by means of the DNN model.

Original data	Transformed data	Output (y)
$Freq =$ wellhead damage frequency value	$\frac{dFreq}{dt} \geq 0$	Risk increase
	$\frac{dFreq}{dt} < 0$	Risk decrease

The simulated indicator values Ind_i , for $i = 1, \dots, 50$, were also transformed into their derivative with respect to time t , in order to define the inputs x_i to the DNN model:

$$x_i = \frac{dInd_i}{dt} \tag{8}$$

Two datasets were created from the overall database:

- Training dataset used to train the DNN model, with 2/3 of the x_i and associated y values (160), and
- Test dataset used to test the DNN model, with about 1/3 of the x_i and associated y values (79).

A code in Python language was written for training and testing. The classifier `tf.contrib.learn.DNNClassifier` from the open-source library TensorFlow (Google LLC, 2018) was used for the DNN model. The DNN model structure (i.e. number of layers and nodes) was inspired by Cheng et al. (2016). Moreover, a multiple linear regression (MLR) model was applied to the same datasets, to provide a term of comparison and evaluate the DNN model ability to predict risk increase.

5. Results

Fig. 8 shows the derivative of risk over time for constant scenario s and consequence c (Eq. (7)) within the considered dataset. For about the first 40 year quarters, the risk value is relatively constant as its derivative oscillates around “0”. Risk variations can be described by the variation of frequency of well damage in Fig. 6, but they are not sudden enough to produce high derivative values. It should be remembered that the frequency of well damage in Fig. 6 is plotted on a logarithmic scale and does not appropriately show the sharp variations of well damage frequency occurring from year quarter 80, which are anyway represented by the risk derivative in Fig. 8.

Fig. 9 shows the results of the risk increase prediction tests by the and the models. The following outcomes are considered:

- true positive (t_p), as correct prediction of risk increase;
- false positive (f_p), as incorrect prediction of risk increase;
- true negative (t_n), as correct prediction of risk decrease; and
- false negative (f_n), as incorrect prediction of risk decrease.

The DNN model has produced fewer false positives and more false negatives than the MLR model. Fig. 10 shows the incorrect predictions over the simulated time. The errors are well distributed along the trend and do not show a specific pattern.

Such results may be also combined to define better representative metrics, as reported by Table 3.

Fewer false positives by the DNN model resulted in higher precision and slightly higher accuracy. However, the higher number of false negatives affected the recall, which is relatively lower than the MLR model.

The results were also evaluated considering a set of tolerance values for the risk derivative. Outcomes obtained for absolute risk derivative lower than specific tolerance values were omitted. Tables 4 and 5 show how respectively the DNN and MLR outcomes gradually change from

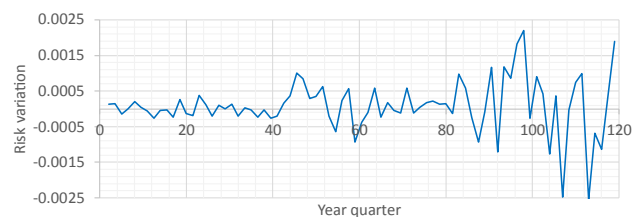


Fig. 8. Derivative of risk over time for constant scenario s and consequence c (Eq. (7)).




Fig. 9. Test results: number of true positives, false positives, true negatives, and false negatives.

the baseline case (null tolerance value) to a tolerance value equal to 0.001, where only 14% of the predictions are considered (the highest peaks in Fig. 8) and no errors are made.

Fig. 11 illustrates the trend of the considered metrics if the tolerance values are varied. The DNN model reports high levels of precision, reaching 100% for a tolerance value equal to 0.0003. Accuracy and recall are also satisfactory, as they reach 100% if the tolerance is equal to 0.0005. On the other hand, the MLR model has higher performance than DNN only in terms of recall, as it reports constantly higher values and reaches 100% for a tolerance value equal to 0.0004. MLR accuracy and precision reach 100% only if tolerance is set to 0.001 due to a persistent false positive error, as seen in Table 5.

6. Discussion

The case study results allow illustrating benefits and limitations of artificial cognition (particularly deep learning) for risk assessment in industry. Having said that, it must be underlined that the main issue is to identify or customize the most suitable model and features given a specific purpose. This requires knowing the state of the art, defining a systematic and evaluation-oriented approach, and applying the right amount of creativity. To this end, the categories of Known/Unknown events (Table 1) and the challenges listed in Section 2.1 are used as a structure to discuss the case study results.

6.1. Known/unknown framework

Paltrinieri et al. (2012b) report an adapted version of the risk management cycle by Merad (2010), which includes the categories of Known/Unknown events (Table 1). Such a framework is used in this work to describe the impact of machine learning on Known/Unknown events (Fig. 12). While machine learning may be considered mostly useful for Known knowns and Unknown knowns, most of the effort is




Table 3
Test results: accuracy, precision and recall.

	DNN model	MLR model	Definitions
Accuracy	83.5%	82.3%	$Acc = \frac{t_p + t_n}{t_p + t_n + f_p + f_n}$
Precision	84.1%	79.6%	$Pr = \frac{t_p}{t_p + f_p}$
Recall	86%	90.7%	$Re = \frac{t_p}{t_p + f_n}$

Table 4
DNN outcomes for specific tolerance values.

Tolerance	Considered predictions (%)	Tp	Tn	Fp	Fn
0	100	37	29	7	6
0.0001	81	33	24	3	4
0.0002	61	25	19	2	2
0.0003	42	21	10	0	2
0.0004	34	17	9	0	1
0.0005	33	17	9	0	0
0.0006	28	13	9	0	0
0.0007	24	12	7	0	0
0.0008	23	11	7	0	0
0.0009	20	9	7	0	0
0.001	14	6	5	0	0

Table 5
MLR outcomes for specific tolerance values.

Tolerance	Considered predictions (%)	Tp	Tn	Fp	Fn
0	100	39	26	10	4
0.0001	81	34	22	5	3
0.0002	61	26	17	4	1
0.0003	42	22	9	1	1
0.0004	34	18	8	1	0
0.0005	33	17	8	1	0
0.0006	28	13	8	1	0
0.0007	24	12	6	1	0
0.0008	23	11	6	1	0
0.0009	20	9	6	1	0
0.001	14	6	5	0	0

required before and in function of acquiring relevant and usable knowledge.

Paltrinieri et al. (2012b) compare an ideal risk management model with the case of an atypical accident (Fig. 12). In this work, we plot the machine learning effort for the ideal case, defined as follows:

$$E = \frac{d(A_i)}{d(K_i)} \tag{12}$$

where E is the machine learning effort equal to the derivative of the awareness A for the unwanted event with respect to the knowledge K of the unwanted event for the ideal case i .

In an initial phase, despite a condition of knowledge and awareness lack, the latter may relatively increase due to reasonable doubt (Merad, 2010). In an ideal case, such reasonable doubt leads to a consolidated

Fig. 10. False positives and false negatives for the prediction of increase in frequency of well-head damage obtained from the models.




Fig. 11. Metrics for specific tolerance values of predictions by (a) the DNN model and (b) the MLR model.

awareness that “something may go wrong” (Kaplan and Garrick, 1981). On the other hand, relative unawareness of a specific accident scenarios and no delayed reasonable doubts can potentially lead to an atypical accident (Paltrinieri et al., 2015, 2011).

The effort in machine learning required by the ideal case is particularly required in the initial phase (phase 1 in Fig. 12). A system for data collection and categorization (the “small things” of section 3) is a necessary support for machine learning, as incomplete and unreliable input data inevitably affect the quality of results. Such a system should be designed at the early stages of risk management for effective implementation of machine learning methods.

The data collection system would also be functional to the realization that there are potential unknown scenarios (Known unknowns). In this phase, new effort should be made to build machine learning models (phase 2 in Fig. 12). The models may already represent a possible response to Known unknowns events, if associated with unsupervised learning (Hastie et al., 2009).

The models are trained (phase 3 in Fig. 12) once relevant knowledge is identified – consciously or unconsciously. In fact, they have more computational power to process all possible variables, so they can detect patterns where human assessors may not be able to see patterns or predictive risk factors.

Once accident scenarios are considered Known knowns, machine learning may help maintain such capability and avoid the potential shift from Known knowns to Known unknowns due to loss of memory (phase 4 in Fig. 12). However, this phase does not require particular effort as the models are supposed to be trained and effective in terms of prediction.

In case of an accident, which may be due to several reasons, such as the presence of an atypical scenario or a loss of memory, a phase of compensation will occur. Such phase represents a response to experiences failure and requires an intense effort for implementing or improving machine learning approaches in the system (phase 5 in Fig. 12).

6.2. Dynamicity

Indicators reporting the system performance on a regular basis represent an opportunity to consider changes and evolutions, and continuously update risk assessment. The example used (Matteini, 2015) simulates the monitoring of 50 indicator categories with regular reading every 6 weeks (Fig. 6). Heterogeneous indicators are considered to describe the safety barrier “stop drive-off”. Considering operational and organizational factors (e.g. number of simulator hours carried out by the DPO in the last three months), in addition to technical ones (e.g. the number of thruster controls failures in the last three months), aims at producing proactive risk evaluation (Paltrinieri et al., 2016a; Scarponi and Paltrinieri, 2016).

Nevertheless, these indicators reflect different projections in time. A technical failure may be directly associated to the accident development, while early operational/organizational deviations have a lower degree of causality and may be disregarded and not registered. Moreover, operational and organizational indicators rely on personnel’s feedback and may be collected less frequently than technical ones. For this reason, sparsity of data may be especially encountered for operational and organizational indicators, and this may undermine the dynamic capabilities of the model.

It must be also mentioned that the DNN model used in this case-study has limitations concerning dynamicity. In fact, every time a new

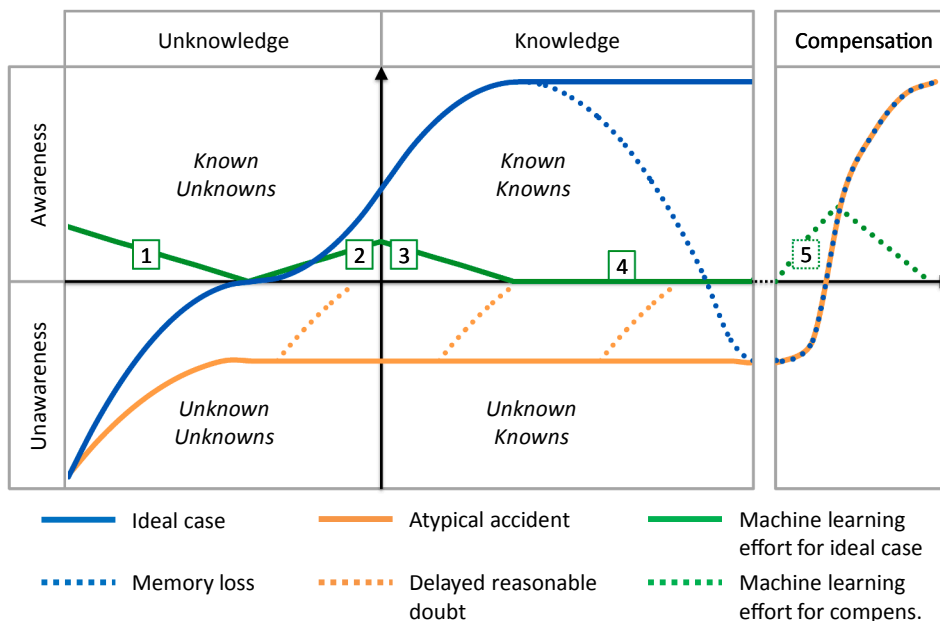


Fig. 12. Risk management cycle of Known/Unknown events through machine learning. Adapted from Paltrinieri et al. (2012b).

set of indicators arrives, the model needs to be re-trained. However, retraining from scratch every time is computationally expensive and delays the time from data arrival to serving an updated model. To tackle this challenge, a warm-starting system is implemented by Cheng et al. (2016), which initializes a new model with embeddings and weights from the previous model.

6.3. Cognition

An artificial cognition model has the potential to capitalize the information collected from indicators and avoid disregard of past lessons. This is made possible by the training sessions, where model features are defined. In this case study, supervised learning was applied: derivatives of the 50 indicator categories were provided together with the associated outputs showing risk increase or decrease. This allows for automatic learning of aggregation structures for input data. Despite the fact that it was not used in this case study, unsupervised learning is also a possibility for machine learning (Hastie et al., 2009). In this case, the desired output is not known (some potential patterns may be anyway provided) and the model aims at drawing inferences in the dataset used.

The additional knowledge dimension for risk definition (as suggested by Aven and Krohn (2014)) is quantitatively represented by the characteristics of the training dataset, such as the number of indicator categories (columns) and values over time (rows), and the number of iterations to minimize the final loss function during model training. In this way, a fundamental concept such as the level of assessment uncertainty can be measured and quantitatively compared.

When we consider such training processes, it is easy to assume that more is better. Nonetheless, as Christian and Griffiths (2016) point out, “the question of how hard to think, and how many factors to consider, is the heart of a knotty problem that statisticians and machine-learning researchers call over-fitting.” The DNN model may have such a sensitivity to input data that the solutions it produces are highly variable. There can be errors in how the data were collected or reported – this is especially true for operational and organizational factors. For instance, collection of the number of DPO delays in the last three months (Fig. 6) depends on DPO’s memory (or honesty) and small mistakes may be amplified in the prediction. For this reason, cross-validating with a test dataset is essential. In this study, a relatively more complex model (DNN) resulted 1.3% more accurate than a linear one (MLR – Table 3), despite the presence of several operational and organizational factors. These factors were simulated to show high volatility (e.g. percentage of time in the last three months with more than an operator monitoring), but we should consider that they may still not be completely realistic.

6.4. Data processing

While machine-learning in general allows overcoming the definition of tangled data aggregation structures and relative weights used for indicators, there are some important differences among the specific techniques. Linear models such as MLR are widely used for prediction purposes. Indicator interactions can be easily memorized through the provided datasets, such as the one in this study (Fig. 8 and Table 3). 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 are suggested for such task (Christian and Griffiths, 2016) and the case study results hint it: when tested with an unknown dataset, the DNN model produced 66 correct predictions of risk increase/decrease against 65 correct predictions by the MLR model (Fig. 8). These results show not only slightly higher accuracy, but also a 5%-higher value for the DNN model precision – compensated by lower recall.

The DNN model seems to perform even better if some tolerance is introduced (Tables 4, 5 and Fig. 11). DNN metrics reach values between

90 and 95% for a value of tolerance equal to about 0.0001 and 81% of predictions are considered, and reach 100% for tolerance equal to 0.0005 with 33% of predictions. For tolerance of 0.0001, MLR accuracy and precision are equal to 87%, while recall is at about 92%. All the MLR metrics reach 100% when tolerance is equal to 0.001 and only 14% of predictions are considered. Such behavior may be explained by the higher sensitivity of DNN models (Christian and Griffiths, 2016), which commit errors only for relatively small risk variations or in the vicinity of stationary points. However, such sensitivity should be appropriately handled as it may lead to over-fitting phenomena (Christian and Griffiths, 2016).

A limitation of DNN is that its 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.

As mentioned, the quality of the model, as with all models, depends on the quality of the data input. For instance, if humans within the system do not think a factor is important, they may not collect the data or include them in the model. In addition, according to the “no free lunch theorem” (Wolpert, 2002), if an algorithm A performs better than algorithm B on a certain problem, it is not necessarily true that A will perform better on other problems. This is why in machine learning it is common to approach the problem by trying more solutions for a particular case. A further model to consider may also be the one suggested by Cheng et al. (2016): a mixed machine-learning model to combine the strengths of both linear and deep approaches. Such technique would allow memorization of registered indicator interactions and generalization of previously unseen ones.

6.5. Emergence

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 better suited to consider rare events, due to their sensitivity to input data and capability to generalize (Cheng et al., 2016).

The case study addressed in this work does not directly address such problems, because it simulates dynamic positioning operations where only deviations from normal conditions and no specific accidents occur. The only relevant result is represented by the demonstration of the potential flexibility of a DNN model. In fact, such a machine learning 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 reminds one of that proposed by Paltrinieri et al. (2013), 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.

Finally, to address the emergence challenge, it is possible to apply 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). For instance, new sets of indicators describing the appropriate operator response to alarms could be introduced in the case study in a second phase without invalidating the evaluation.

6.6. Usability

The case study showed how a machine approach allows predicting the overall risk of well damage increase or decrease based on the variation of singular technical, operational and organizational indicators. This approach may be used for both real-time risk assessment of the overall system, and simulation of possible future scenarios. Understanding whether the system is, or may be, improving or worsening in terms of risk is a fundamental support to safety-related decision-making. In fact, risk informed decisions are used in a number of circumstances where something of value is at stake (Kongsvik et al., 2015).

The metrics used to assess the performance of DNN and MLR models may also inform the model suitability for specific decision-making tasks. In fact, in addition to accuracy, prediction and recall should be considered. The former shows the ratio of correct risk increase predictions over all the risk increase predictions by the model. The latter shows the ratio of correct risk increase predictions over all the real risk increase events. In this case, the model predicts risk increase or decrease following normal operations. For this reason, both risk increase and decrease are important and relatively frequent, and what we should search for is model accuracy, and subsequently precision. Given the relatively low criticality of the prediction target, the results obtained in this application may be considered acceptable. Moreover, results can be further improved if tolerance is set (Tables 4, 5 and Fig. 11).

In case of predictions of rare events, such as major accidents, recall assumes a primary role. For instance, the highest blowout probability during offshore drilling operations is estimated by Khakzad et al. (2013b) as 0.00002. A model never predicting any blowout would have accuracy next to 100%, but precision and recall equal to 0%. Moreover, the criticality of such accidents tolerates conservative predictions. For this reason, recall, which disregards false positives and focuses on true positives, would be the metric to prioritize.

Further processing of case study indicators could also have led to the prediction of a risk index value, but this was not carried out for the sake of brevity. Prediction of a risk index value would have enabled the use of a risk barometer such as the one depicted in Fig. 7. Such risk visualization tool (Edwin et al., 2016) may be used to communicate risk predicted with a traffic-light colour code. Risk communication is an important purpose of risk assessment and essential to raise general awareness. Effective risk communication should be carried out among the main parties of an industrial site (Paltrinieri et al., 2012b). In fact, participation by multiple parties in information sharing amplifies its benefits, especially when the parties face common risks (Phimister et al., 2004).

7. Conclusions

Through this work, we have addressed what we believe are the main current challenges of industrial risk assessment and we have suggested an approach based on machine learning as a possible answer. A DNN model has been used for the risk assessment of a drive-off scenario involving an Oil & Gas drilling rig. The developed model aims to predict risk increase or decrease as the system conditions (described by performance indicators) change. Results from a test on the DNN model and a comparison with a MLR model show that the former is more accurate for dynamic assessment and presents the required flexibility to deal with unexpected events. Nevertheless, despite apparent suitability to (partially) overcome risk assessment challenges, intrinsic DNN limitations should always be taken into account. For instance, its high model sensitivity does not tolerate inaccurate indicators and can potentially lead to over-fitting. 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.

Acknowledgements

This research was supported by the project Lo-Risk (“Learning about Risk”), supported by the Norwegian University of Science and Technology – NTNU (Onsager fellowship).

References

- Apostolakis, G.E., 2004. How useful is quantitative risk assessment? *Risk Anal.* 24, 515–520. <https://doi.org/10.1111/j.0272-4332.2004.00455.x>.
- Aven, T., 2012. The risk concept—historical and recent development trends. *Reliab. Eng. Syst. Saf.* 99, 33–44. <https://doi.org/10.1016/j.res.2011.11.006>.
- Aven, T., Krohn, B.S., 2014. A new perspective on how to understand, assess and manage risk and the unforeseen. *Reliab. Eng. Syst. Saf.* 121, 1–10.
- Bucelli, M., Paltrinieri, N., Landucci, G., 2017. Integrated risk assessment for oil and gas installations in sensitive areas. *Ocean Eng.* <https://doi.org/10.1016/j.oceaneng.2017.12.035>.
- Center for Chemical Process Safety, 2000. Guidelines for Chemical Process Quantitative Risk Analysis. American Institute of Chemical Engineers (AIChE), New York, USA.
- CGE Risk Management Solutions B.V., 2016. BowTieXP [WWW Document].
- Chen, H., Moan, T., Verhoeven, H., 2008. Safety of dynamic positioning operations on mobile offshore drilling units. *Reliab. Eng. Syst. Saf.* 93, 1072–1090. <https://doi.org/10.1016/J.RESS.2007.04.003>.
- Cheng, H.-T., Koc, L., Harmsen, J., Shaked, T., Chandra, T., Aradhya, 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.
- Comfort, L.K., 2019. The Dynamics of Risk: Changing Technologies, Complex Systems, and Collective Action in Seismic Policy. Princeton University Press, Princeton.
- Comfort, L.K., 1999. Shared Risk: Complex Systems in Seismic Policy. Pergamon, Amsterdam and Oxford.
- Creeedy, G.D., 2011. Quantitative risk assessment: how realistic are those frequency assumptions? *J. Loss Prev. Process Ind.* 24, 203–207. <https://doi.org/10.1016/j.jlpp.2010.08.013>.
- Sornette, D., 2009. Dragon-Kings, Black Swans and the Prediction of Crises. ETH Zurich, Chair of Systems Design.
- De Marchi, B., Ravetz, J.R., 1999. Risk management and governance: a post-normal science approach. *Futures* 31, 743–757. [https://doi.org/10.1016/S0016-3287\(99\)00030-0](https://doi.org/10.1016/S0016-3287(99)00030-0).
- Durga Rao, K., Gopika, V., Sanyasi Rao, V.V.S., Kushwaha, H.S., Verma, A.K., Srividya, A., 2009. Dynamic fault tree analysis using Monte Carlo simulation in probabilistic safety assessment. *Reliab. Eng. Syst. Saf.* 94, 872–883. <https://doi.org/10.1016/j.res.2008.09.007>.
- Diekmann, E.J., 1992. Risk analysis: lessons from artificial intelligence. *Int. J. Proj. Manag.* 10, 75–80. [10.1016/0263-7863\(92\)90059-1](https://doi.org/10.1016/0263-7863(92)90059-1).
- Edwin, N.J., Paltrinieri, N., Østerlie, T., 2016. Risk Metrics and Dynamic Risk Visualization. In: Dynamic Risk Analysis in the Chemical and Petroleum Industry: Evolution and Interaction with Parallel Disciplines in the Perspective of Industrial Application. [10.1016/B978-0-12-803765-2.00013-5](https://doi.org/10.1016/B978-0-12-803765-2.00013-5).
- Etterlid, D.E., Etterlid, D., 2013. iSee – Visualization of risk related factors. IO Conf.
- 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.
- European Parliament and Council, 1982. Council Directive 82/501/EEC of 24 June 1982 on the major-accident hazards of certain industrial activities. *Off. J. Eur. Union* 1–18.
- Giddings, I.C., 2013. IMO Guidelines for vessels with dynamic positioning systems. Dynamic Positioning Conference, Houston, Texas, U.S.
- 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. < www.tensorflow.org > .
- Hansen, H.N., 2015. Goliath Barrier Management, in: Barrierstyring i Praksis. ESRA, Oslo, Norway.
- Hastie, T., Tibshirani, R., Friedman, J., 2009. Unsupervised learning. In: The Elements of Statistical Learning. Springer, pp. 485–585.
- Hauge, S., Okstad, E., Paltrinieri, N., Edwin, N., Vatn, J., Bodsberg, L., 2015. Handbook for monitoring of barrier status and associated risk in the operational phase. SINTEF F27045. Center for Integrated Operations in the Petroleum Industry, Trondheim, Norway, Norway.
- Haugen, S., Vinnem, J.E., 2015. Perspectives on risk and the unforeseen. *Reliab. Eng. Syst. Saf.* 137, 1–5.
- Hutchins, E., 1995. Cognition in the Wild. MIT press.
- Kaplan, S., Garrick, B.J., 1981. On the quantitative definition of risk. *Risk Anal.* 1, 11–27. <https://doi.org/10.1111/j.1539-6924.1981.tb01350.x>.
- Kaufmann, S.A., 1993. The origins of order: self-organization and selection in evolution.
- Khakzad, N., 2015. Application of dynamic Bayesian network to risk analysis of domino effects in chemical infrastructures. *Reliab. Eng. Syst. Saf.* 138, 263–272. <https://doi.org/10.1016/j.res.2015.02.007>.
- Khakzad, N., Khan, F., Amyotte, P., 2013a. Risk-based design of process systems using discrete-time Bayesian networks. *Reliab. Eng. Syst. Saf.* 109, 5–17. <https://doi.org/>

- 10.1016/j.res.2012.07.009.
- Khakzad, N., Khan, F., Amyotte, P., 2013b. Quantitative risk analysis of offshore drilling operations: a Bayesian approach. *Saf. Sci.* 57, 108–117. <https://doi.org/10.1016/J.SSCI.2013.01.022>.
- Khakzad, N., Khan, F., Paltrinieri, N., 2014. On the application of near accident data to risk analysis of major accidents. *Reliab. Eng. Syst. Saf.* 126. <https://doi.org/10.1016/j.res.2014.01.015>.
- Khan, F., Hashemi, S.J., Paltrinieri, N., Amyotte, P., Cozzani, V., Reniers, G., 2016. Dynamic risk management: a contemporary approach to process safety management. *Curr. Opin. Chem. Eng.* 14. <https://doi.org/10.1016/j.coche.2016.07.006>.
- King, G., Zeng, L., 2001. Logistic regression in rare events data. *Polit. Anal.* 9, 137–163.
- Kongsvik, T., Almklov, P., Haavik, T., Haugen, S., Vinnem, J.E., Schiefloe, P.M., 2015. Decisions and decision support for major accident prevention in the process industries. *J. Loss Prev. Process Ind.* 35, 85–94. <https://doi.org/10.1016/j.jlp.2015.03.018>.
- Landucci, Gabriele, Paltrinieri, N., 2016a. A methodology for frequency tailorization dedicated to the Oil & Gas sector. *Process Saf. Environ. Prot.* 104, 123–141. <https://doi.org/10.1016/j.psep.2016.08.012>.
- Landucci, G., Paltrinieri, N., 2016b. Dynamic evaluation of risk: from safety indicators to proactive techniques. *Chem. Eng. Trans.* <https://doi.org/10.3303/CET1653029>.
- Landucci, G., Pontiggia, M., Paltrinieri, N., Cozzani, V., 2016a. Dynamic Consequence Analysis through Computational Fluid Dynamics Modeling. In: *Dynamic Risk Analysis in the Chemical and Petroleum Industry: Evolution and Interaction with Parallel Disciplines in the Perspective of Industrial Application*. 10.1016/B978-0-12-803765-2.00009-3.
- Landucci, G., Pontiggia, M., Paltrinieri, N., Cozzani, V., 2016b. Computational Fluid Dynamics Modeling: 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*. 10.1016/B978-0-12-803765-2.00010-X.
- LaPorte, T.R., Consolini, P.M., 1991. Working in practice but not in theory: theoretical challenges of high-reliability organizations. *J. Public Adm. Res. Theory J-PART 1*, 19–48.
- Lasi, H., Fettke, P., Kemper, H.-G., Feld, T., Hoffmann, M., 2014. *Industry 4.0*. *Bus. Inf. Syst. Eng.* 6, 239–242.
- Matteini, A., 2015. Human factors and dynamic risk analysis: a case-study in Oil&Gas drilling. Bologna, Italy.
- Merad, M., 2010. *Aide à la décision et expertise en gestion des risques*. Lavoisier, Paris.
- Musgrave, G.L., 2013. James Owen Weatherall, *The Physics of Wall Street: A Brief History of Predicting the Unpredictable*. *Bus. Econ.* 48, 203–204.
- Nivolianitou, Z.S., Leopoulou, V.N., Konstantinidou, M., 2004. Comparison of techniques for accident scenario analysis in hazardous systems. *J. Loss Prev. Process Ind.* 17, 467–475. <https://doi.org/10.1016/j.jlp.2004.08.001>.
- Nobre, F.S., 2009. Designing Future Information Management Systems. IGI Global.
- Noh, Y., Chang, K., Seo, Y., Chang, D., 2014. Risk-based determination of design pressure of LNG fuel storage tanks based on dynamic process simulation combined with Monte Carlo method. *Reliab. Eng. Syst. Saf.* 129, 76–82. <https://doi.org/10.1016/j.res.2014.04.018>.
- Nývlt, O., Haugen, S., Ferkl, L., 2015. Complex accident scenarios modelled and analysed by Stochastic Petri Nets. *Reliab. Eng. Syst. Saf.* 142, 539–555. <https://doi.org/10.1016/j.res.2015.06.015>.
- Nývlt, O., Rausand, M., 2012. Dependencies in event trees analyzed by Petri nets. *Reliab. Eng. Syst. Saf.* 104, 45–57. <https://doi.org/10.1016/j.res.2012.03.013>.
- Øien, K., Utne, I.B., Herrera, I.A., 2011. Building safety indicators: Part 1 - theoretical foundation. *Saf. Sci.* 49, 148–161. <https://doi.org/10.1016/j.ssci.2010.05.012>.
- Paltrinieri, N., Dechy, N., Salzano, E., Wardman, M., Cozzani, V., 2012a. Lessons learned from Toulouse and Buncefield Disasters: from risk analysis failures to the identification of atypical scenarios through a better knowledge management. *Risk Anal.* 32. <https://doi.org/10.1111/j.1539-6924.2011.01749.x>.
- Paltrinieri, Nicola, Khan, F., 2016. Dynamic Risk Analysis in the Chemical and Petroleum Industry, *Dynamic Risk Analysis in the Chemical and Petroleum Industry: Evolution and Interaction with Parallel Disciplines in the Perspective of Industrial Application*. Butterworth-Heinemann. 10.1016/B978-0-12-803765-2.01001-5.
- 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*. 10.1016/B978-0-12-803765-2.00002-0.
- Paltrinieri, N., Khan, F., Amyotte, P., Cozzani, V., 2014. Dynamic approach to risk management: application to the Hoggans metal dust accidents. *Process Saf. Environ. Prot.* 92. <https://doi.org/10.1016/j.psep.2013.11.008>.
- Paltrinieri, N., Khan, F., Cozzani, V., 2015. Coupling of advanced techniques for dynamic risk management. *J. Risk Res.* 18, 910–930. <https://doi.org/10.1080/13669877.2014.919515>.
- Paltrinieri, N., Landucci, G., Nelson, W.R., Hauge, S., 2016a. Proactive Approaches of Dynamic Risk Assessment Based on Indicators. In: *Dynamic Risk Analysis in the Chemical and Petroleum Industry: Evolution and Interaction with Parallel Disciplines in the Perspective of Industrial Application*. 10.1016/B978-0-12-803765-2.00006-8.
- Paltrinieri, N., Massaiu, S., Matteini, A., 2016b. Human Reliability Analysis in the Petroleum Industry: Tutorial and Examples. Butterworth-Heinemann, pp. 181–192. 10.1016/B978-0-12-803765-2.00015-9.
- Paltrinieri, N., Øien, K., Cozzani, V., 2012b. Assessment and comparison of two early warning indicator methods in the perspective of prevention of atypical accident scenarios. *Reliab. Eng. Syst. Saf.* 108. <https://doi.org/10.1016/j.res.2012.06.017>.
- 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., Buston, J., Wardman, M., Cozzani, V., 2013. Dynamic procedure for atypical scenarios identification (DyPASI): a new systematic HAZID tool. *J. Loss Prev. Process Ind.* 26. <https://doi.org/10.1016/j.jlp.2013.01.006>.
- Pasman, H., Reniers, G., 2014. Past, present and future of Quantitative Risk Assessment (QRA) and the incentive it obtained from Land-Use Planning (LUP). *J. Loss Prev. Process Ind.* 28, 2–9.
- 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>.
- Petroleum Safety Authority, 2013. Principles for barrier management in the petroleum industry. PSA, Stavanger, Norway.
- Phimister, J.R., Bier, V.M., Kunreuther, H.C., Engineering, N.A., 2004. *Accident Precursor Analysis and Management: Reducing Technological Risk Through Diligence*. National Academies Press.
- Prigogine, I., Stengers, I., 1984. *Order Out of Chaos: Man's New Dialogue with Nature*. Bantam Books.
- PSA, 2016. Trends in risk level in the petroleum activity (RNNP) [WWW Document]. < <http://www.psa.no/about-rnnp/category911.html> > .
- Rochlin, G.L., La Porte, T.R., Roberts, K.H., 1987. The self-designing high-reliability organization: aircraft carrier flight operations at sea. *Nav. War Coll. Rev.* 40, 76–90.
- Roe, E., Schulman, P.R., 2008. *High Reliability Management: Operating on the Edge*. Stanford University Press.
- 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>.
- Scarpioni, G.E., Paltrinieri, N., 2016. Comparison and Complementarity between Reactive and Proactive Approaches. In: *Dynamic Risk Analysis in the Chemical and Petroleum Industry: Evolution and Interaction with Parallel Disciplines in the Perspective of Industrial Application*. 10.1016/B978-0-12-803765-2.00008-1.
- Schumacher, A., Erol, S., Sihni, W., 2016. A maturity model for assessing industry 4.0 readiness and maturity of manufacturing enterprises. *Proc. CIRP* 52, 161–166. <https://doi.org/10.1016/J.PROCIR.2016.07.040>.
- Statoil, 2013. Annual report Reporting Compendium. Annual report 2012. Statoil, Oslo, Norway.
- 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](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.
- Targoutzidis, A., 2012. A Monte Carlo simulation for the assessment of Bayesian updating in dynamic systems. *Reliab. Eng. Syst. Saf.* 100, 125–132. <https://doi.org/10.1016/j.res.2011.12.020>.
- Venkatesan, R., Er, M.J., 2016. A novel progressive learning technique for multi-class classification. *Neurocomputing* 207, 310–321.
- Villa, V., Paltrinieri, N., Khan, F., Cozzani, V., 2016a. Towards dynamic risk analysis: a review of the risk assessment approach and its limitations in the chemical process industry. *Saf. Sci.* 89. <https://doi.org/10.1016/j.ssci.2016.06.002>.
- Villa, V., Paltrinieri, N., Khan, F., Cozzani, V., 2016b. A Short Overview of Risk Analysis Background and Recent Developments. In: *Dynamic Risk Analysis in the Chemical and Petroleum Industry: Evolution and Interaction with Parallel Disciplines in the Perspective of Industrial Application*. 10.1016/B978-0-12-803765-2.00001-9.
- Wang, K.S., Hsu, F.S., Liu, P.P., 2002. Modeling the bathtub shape hazard rate function in terms of reliability. *Reliab. Eng. Syst. Saf.* 75, 397–406. [https://doi.org/10.1016/S0951-8320\(01\)00124-7](https://doi.org/10.1016/S0951-8320(01)00124-7).
- Wang, S., Li, D., Zhang, C., 2016. Towards smart factory for industry 4.0: a self-organized multi-agent system with big data based feedback and coordination. *Comput. Networks* 101, 158–168. <https://doi.org/10.1016/J.COMNET.2015.12.017>.
- Weick, K.E., Roberts, K.H., 1993. *Collective mind in organizations: Heedful interrelating on flight decks*. *Adm. Sci. Q.* 357–381.
- Weick, K.E., Sutcliffe, K.M., 2001. *Managing the Unexpected: Assuring High Performance in an Age of Complexity*, published by JosseyBass, a John Wiley & Sons. Inc. Co.
- Wolpert, D.H., 2002. The supervised learning no-free-lunch theorems. *Soft Computing and Industry* 25–42.
- Yang, X., Haugen, S., 2015. Classification of risk to support decision-making in hazardous processes. *Saf. Sci.* 80, 115–126. <https://doi.org/10.1016/J.SSCI.2015.07.011>.
- Yang, X., Haugen, S., Paltrinieri, N., 2017. Clarifying the concept of operational risk assessment in the oil and gas industry. *Saf. Sci.* 10.1016/j.ssci.2017.12.019.
- Zhou, J., Reniers, G., 2017. Petri-net based cascading effect analysis of vapor cloud explosions. *J. Loss Prev. Process Ind.* 48, 118–125. <https://doi.org/10.1016/J.JLP.2017.04.017>.
- Zhou, J., Reniers, G., 2016a. Petri-net based simulation analysis for emergency response to multiple simultaneous large-scale fires. *J. Loss Prev. Process Ind.* 40, 554–562. <https://doi.org/10.1016/J.JLP.2016.01.026>.
- Zhou, J., Reniers, G., 2016b. Petri-net based modeling and queuing analysis for resource-oriented cooperation of emergency response actions. *Process Saf. Environ. Prot.* 102, 567–576. <https://doi.org/10.1016/J.PSEP.2016.05.013>.
- Zhou, J., Reniers, G., Zhang, L., 2017. A weighted fuzzy Petri-net based approach for security risk assessment in the chemical industry. *Chem. Eng. Sci.* 174, 136–145. <https://doi.org/10.1016/J.CES.2017.09.002>.