CHEMICAL ENGINEERING TRANSACTIONS

VOL. 82, 2020

A publication of AIDIC
The Italian Association of Chemical Engineering
Online at www.cetjournal.it

Guest Editors: Bruno Fabiano, Valerio Cozzani, Genserik Reniers

ISBN 978-88-95608-80-8; **ISSN** 2283-9216

DOI: 10.3303/CET2082029

Meta-learning for Safety Management

Nicola Paltrinieri[a,*], Riccardo Patriarca[b], Elena Stefana[c], Francisco Brocal[d], Genserik Reniers[e]

[a]Department of Mechanical and Industrial Engineering, NTNU, Trondheim, Norway
[b]Department of Mechanical and Aerospace Engineering, Sapienza University of Rome, Italy
[c]Department of Mechanical and Industrial Engineering, University of Brescia, Italy
[d]Department of Physics, Systems Engineering and Signal Theory, University of Alicante, Spain
[e]Faculty of Applied Economics, University of Antwerp Operations Research Group, Belgium
nicola.paltrinieri@ntnu.no

The experience gathered from normal industrial operations allows us to associate its degrading conditions with the potential for an accident. Such association is the basis for the definition of the system risk and appropriate safety measures. If a skilled operator observes further degrading conditions, his/her mind quickly learns from this new experience, derives an updated risk level, and tunes the safety measures. Similarly, safety management techniques aim to construct a risk model while learning from past and new observations with the purpose to warn of an imminent accident. However, the model can be tested only in hindsight, after the occurrence (or the missed occurrence) of an accident. How can we generalise and model the risk analysis learning process? How can we optimise its configuration towards new observations? This study discusses these issues of meta-learning for safety management by considering the case study of a drive-off scenario involving an oil and gas drilling rig, for which a risk assessment approach based on machine learning is developed. The results indicate the way forward for a generalisation of risk analysis learning processes and their optimisation.

1. Introduction

Shifts in our assessment of risk are continuously imposed by emergence of new knowledge, reshaping the limits of our actions. This is particularly important in high-risk technical sectors, striving for enhanced system performance, but where accidents can affect many people. A classic definition of risk is given by Kaplan and Garrick (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):

$$R = f(s, p, c) \tag{1}$$

The continuous occurrence of major accidents resulting from the failure to learn from experience are reminders of the details that cannot be framed by Eq(1) (Paltrinieri et al., 2012). Numerous attempts have been made by analysts and scholars to capture the notion of risk in a more meaningful way. Aven (Aven, 2012) provides a thorough review of risk definitions, while Villa et al. (Villa et al., 2016a, 2016b) show that differences in risk definition affect the approach adopted for its assessment and management. Aven and Krohn (Aven and Krohn, 2014) suggest including also the knowledge dimension in the definition of risk, as the accumulated knowledge is an intrinsic feature of the assessment. Instead, the standard ISO 31000 defines risk as the effect of uncertainty on objectives (ISO, 2018). This gives important insight on how we should treat risk analysis results and promotes continuous improvement of the analysis itself – we become aware of how uncertainty is an inescapable companion and that we should cope with it (De Marchi and Ravetz, 1999).
Even if we can assess risk with all available knowledge, we would provide a risk picture that is “frozen” in time, while the system is changing around it. The conditions considered in time 0 may not be valid anymore in time *n*. Calibration and correction based on new evidence would possibly allow risk analysis to consider evolving conditions and reflect reality and its results. Such dynamic approach to risk management is theorized and reviewed by several previous works (Bubbico et al., 2020; Khan et al., 2016; Lee et al., 2019; Paltrinieri et al.,

Paper Received: 30 January 2020; Revised: 27 April 2020; Accepted: 31 July 2020
Please cite this article as: Paltrinieri N., Patriarca R., Stefana E., Brocal F., Reniers G., 2020, Meta-learning for Safety Management, Chemical Engineering Transactions, 82, 169-174 DOI:10.3303/CET2082029

2014). However, Paltrinieri et al. (Paltrinieri et al., 2019) highlight a set of overall challenges that are still present within the field of risk analysis despite the most recent progress. In particular, they focus on cognition and emergence. They wonder how we can learn from relevant lessons to improve risk analysis. Unwanted events and experts can provide valuable insight. Capitalising such knowledge in a systematic way would prevent accident repetition. They wonder how we can 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, or lack of risk awareness.

1.1 Meta-learning

In industrial sectors where the sense of risk is constantly present, such as oil and gas, experience gathered from operating a technical system allows skilled operators to associate the system conditions with a specific level of accident risk (Duan, 2018; Hailwood, 2016). For example, corrosion on a vessel may eventually lead to its catastrophic rupture and its presence would be associated with a relatively high risk. This experience allows assessing the risk whenever we find the same conditions of corrosion. Instead, if the conditions are only similar (other corrosion mechanisms) or new (mechanical fatigue), the operators' mind will respectively derive an adequate risk level by quickly learning from this new experience. The system conditions are features that may be reported in a vector X, the risk is the target variable R, and the experience is the dataset D of the operators' observations. D is the basis used by the operators' mind to build a model $f(X) = R$. The field of Risk Analysis aims to provide an artificial risk model $f_\theta(X)$. The model has a structure configured by a set of parameters $\theta = \{\phi_1, \dots, \phi_n\}$, which are defined (trained) on the historical observations collected in D. An observation batch from D can allow us to test the model and estimate its risk prediction performance $RPP\big(R, f(X)\big)$, i.e. the capacity to warn of a potential accident. However, RPP is unknown if the model is required to process system conditions X_{T+1} at a time $T+1$ that were not observed before. In fact, the model can be tested only in hindsight, after the occurrence (or the missed occurrence) of an accident. This translates the fundamental challenges of cognition and emergence as follows. Cognition: the first challenge addresses the risk analysis learning process $L(\theta)$ and the expected predictive performance associated with a configuration θ for a given dataset D.

$$L(\theta) = E\Big(RPP\big(R, f(X,\theta)\big)\Big|D\Big) \tag{2}$$

Emergence: the second challenge addresses the configuration θ^* for the best learning process $L(\theta^*)$ on a distribution of datasets, including potentially unseen datasets at time $t = T+1$.

$$\theta^* = \underset{\theta}{\operatorname{argmax}}\, E\Big(RPP\big(R, f(X,\theta)\big)\Big|D_t\Big), \qquad t = 1, \dots, T+1 \tag{3}$$

This study addresses these issues by considering the case study of a drive-off scenario involving an oil and gas drilling rig, for which a risk assessment approach based on machine learning is developed. Through the case-study, we discuss the risk analysis learning model and its optimisation towards new observations, in order to apprehend the emergence of unknown risks.

2. Drive-off scenario involving an oil and gas drilling rig

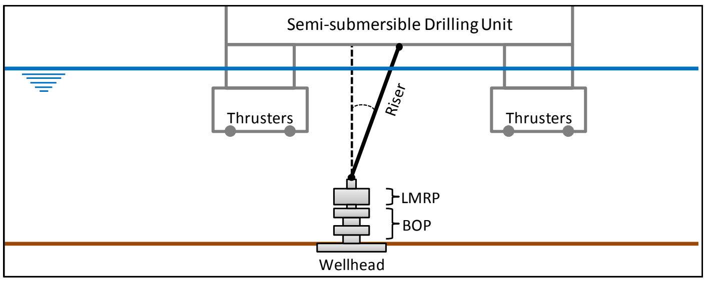

Figure 1: Position of a semi-submersible drilling unit above the wellhead

In order to avoid potential damage during drilling operations for a new offshore Oil and gas well, a semi-submersible drilling unit should maintain the position above the wellhead (Figure 1). 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 several 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 initiate 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). A number of works (Matteini, 2015; Paltrinieri et al., 2019, 2016) address the details of occurrence and development of drive-off scenarios. Relevant indicators are defined to assess the performance of safety barriers and related systems. Examples of these indicators are the following.

- Thruster control failures in the last three months.
- Thruster monitoring sensors failures in the last three months.
- Simulator hours carried out by the DPO in the last three months.
- Inadequate DPO communication events in the last three months.
- Delays in DPO shifts in the last three months.
- Percentage of time in the last three months with more than one operator monitoring.

The simulations of their trends for a period of 30 years can be found in the literature (Paltrinieri et al., 2019). They are inspired to the typical bathtub curve for technical elements (Wang et al., 2002) and relevant expert judgment for the remaining elements. As shown by Bucelli et al. (Bucelli et al., 2017), indicator values (representing the system conditions X) 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. Outcome frequencies are an expression of the scenario probability *p* mentioned in Eq(1) and, in turn, of the risk *R*. If we assume that the other factors are constant, this represents a simplified model $f(X) = R$. However, Matteini (Matteini, 2015) points out a certain complexity within the hierarchical barrier model, which may be due to a tangled structure and an unclear approach to assign relative weights to single model elements. For this reason, a machine learning approach bypassing the construction of such hierarchies and aggregation rules is suggested.

3. Method

Machine learning refers to techniques aiming to program computers to learn from experience (Samuel, 1959). It allows computational models to learn representations of data with certain levels of abstraction. A computer may be trained to assess risk for safety-critical industries such as oil and gas through machine learning techniques. A large amount of information in the form of the mentioned indicators may be used for training. Once the model has learned risk categorisation and created an artificial risk model $f_\theta(X)$, it uses its knowledge to assess real-time risk from the state of the monitored system, e.g. an offshore oil and gas drilling rig. The machine learning technique used for this study is the Multiple Linear Regression (MLR) (Bottenberg and Ward, 1963), where the variables are the indicators. 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:

$$\frac{dFreq}{dt} \approx \frac{dR}{dt} \tag{4}$$

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 1). The simulated indicator values Ind were also transformed into their derivative with respect to time *t*, in order to define the inputs X to the model $f_\theta(X)$:

$$X = \frac{dInd}{dt} \tag{5}$$

Two datasets were created:

- training dataset with 2/3 of the X and associated R values (160), and
- test dataset used to test the model f_θ, with about 1/3 of the X and associated R values (79).

A code in Python language was written for training and testing. The classifier *tf.contrib.learn.LinearClassifier* from the open-source library TensorFlow (Google LLC, 2018) was used for the model.

Table 1: Definition of the output used as risk index. Adapted from (Paltrinieri et al., 2019)

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

4. Results

The main results given by the study is the creation of an MLR model $f_\theta(X)$ predicting increase of wellhead damage risk given the indicator trends for scenario of a drilling rig drive-off. The model test has the purpose to estimate the prediction performance $RPP(R, f(X))$ based on a known test dataset. The model elaborates a risk increase probability value for each dataset record (Figure 2).

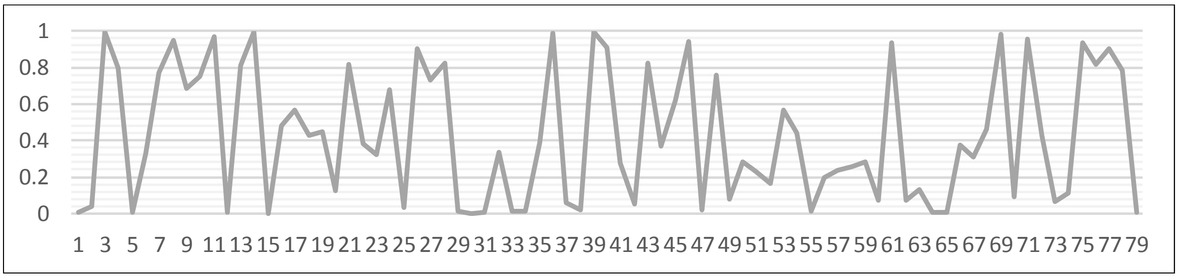

Figure 2: Risk increase probability values for the test dataset

A decision on the risk increase prediction is made by the model by means of a default probability threshold equal to 0.5; meaning that the model predicts risk increase for probability values higher than 0.5. Figure 3 shows the results of the risk increase prediction tests. The following outcomes are considered: i) true positive (t_p), as correct prediction of risk increase; ii) false positive (f_p), as incorrect prediction of risk increase; iii) true negative (t_n), as correct prediction of risk decrease; and iv) false negative (f_n), as incorrect prediction of risk decrease. Only 4 cases are predicted as risk decrease while the risk is actually increasing. However, the model wrongly predicts 10 risk increases while the risk is decreasing. This is also reflected by the metrics considered.

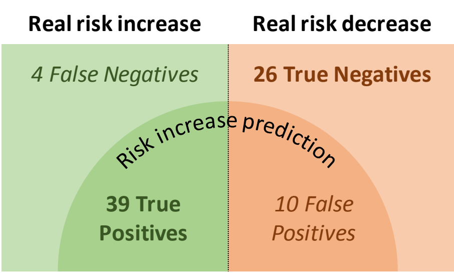

Metric	Result	Definition
Accuracy	82.3%	$Acc = (t_p + t_n)/(t_p + t_n + f_p + f_n)$
Precision	79.6%	$Pr = t_p/(t_p + f_p)$
Recall	90.7%	$Re = t_p/(t_p + f_n)$

Figure 3: Test results: number of true positives, false positives, true negatives, false negatives, and related metrics (given a threshold=0.5)

5. Discussion

The results provide a model $f_\theta(X)$ and an estimation of its risk prediction performance $RPP(R, f(X))$. However, the aspect of meta-learning is not directly addressed as long as the cognition and emergence challenges are out of the picture. To tackle the cognition challenge, the parameter θ is identified in the model decision threshold. Figure 4 shows a PR (precision recall) curve obtained by variating the threshold. A number of other configuration parameters affecting the expected predictive performance (Goodfellow et al., 2016) may be also considered, but they are out of the scope of this study. Precision and recall are an important measure of the predictive performance and have intrinsic differences. The former shows the ratio of correct risk increase predictions over all the risk increase predictions by the model, while the latter the ratio of correct risk increase predictions over all the real risk increase events. In this case, the model predicts the increase or decrease of wellhead damage risk due to drive-off following normal drilling operations. For this reason, both risk increase and decrease are relatively frequent and none of them prevails on the other. Given the relatively low criticality of the prediction target, the model performance should be improved based on accuracy and precision. On the other hand, for unbalanced dataset, i.e. in case of predictions of rare events such as major accidents, recall assumes a primary role (Paltrinieri et al., 2019). To address the emergence challenge, we must ensure the optimisation of the learning process and enhance the predictive performance for unseen conditions. For this reason, the search for the highest F-measure may be integrated into the learning process. This method would allow defining the configuration θ^* for which the performance in terms of either precision or recall is enhanced. The F-measure is defined as follows (Sasaki, 2007):

$$F_\beta = (\beta^2 + 1) \cdot \frac{Precision \cdot Recall}{\beta^2 \cdot Precision + Recall} \tag{6}$$

β is used to weight the variables: when $\beta > 1$, the F-measure is more recall-oriented. When $\beta < 1$, it becomes more precision-oriented. When $\beta = 1$, the F-measure represents the harmonic mean between precision and recall. Figure 4 shows F-measures for β equal to 0.5, 1, and 1.5. Considered that in this case study we search for model accuracy and precision, the most appropriate threshold is equal to 0.6, as it maximises both the former and a precision-oriented F-measure such as $F_{1.5}$.

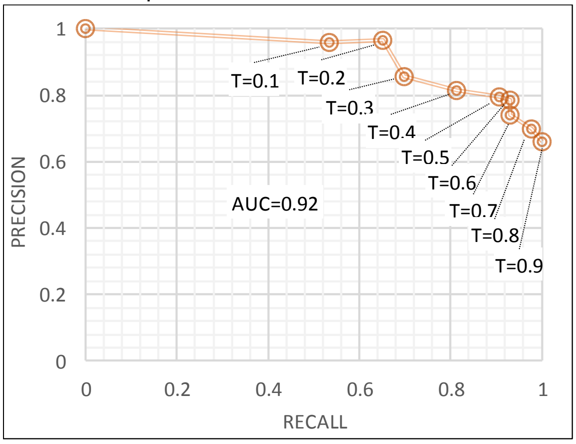

Threshold	Accuracy	$F_{0.5}$	F_1	$F_{1.5}$
0.1	0.734177	0.827338	0.686567	0.619048
0.2	0.797468	0.880503	0.777778	0.723658
0.3	0.772152	0.819672	0.769231	0.740038
0.4	0.797468	0.813953	0.813953	0.813953
0.5	0.822785	0.8159	0.847826	0.86964
0.6	0.822785	0.809717	0.851064	0.879865
0.7	0.78481	0.772201	0.824742	0.862355
0.8	0.759494	0.742049	0.815534	0.870813
0.9	0.721519	0.709571	0.796296	0.863988

Figure 4: PR (precision recall) curve variating the threshold (T) (AUC stands for area under the curve) and related accuracy and F-measures for β equal to 0.5, 1, and 1.5 (highest values highlighted in yellow)

6. Conclusions

This contribution illustrates the preliminary results of a meta-learning study for a drilling rig safety management. A drive-off scenario is considered, and the increase of wellhead damage risk is predicted. An attempt to generalise and model the risk analysis learning process is made, but an actual formalisation is still missing. However, the decision threshold is identified as a configuration parameter to optimise. This would allow improving the model performance towards new observations. The search for the highest F-measure is suggested as an integration to the actual learning process. The F-measure should promote either precision or recall based on the event that is being predicted. The former should be considered for low-criticality events such as the risk increase considered in this study. An optimised threshold for the considered case study is obtained through this approach with the purpose of demonstrating its efficiency. Furthermore, future research efforts may be devoted at exploring the validity of the indicators currently used for the analysis, encompassing recent research grounded in resilience management and normal work operations (Patriarca et al., 2019, 2018). These results are only a first step into the domain of meta-learning for safety management but indicate the way forward for a generalisation of risk analysis learning processes and their optimisation.

Acknowledgments

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

Aven T., 2012, The risk concept—historical and recent development trends, Reliab. Eng. Syst. Saf. 99, 33–44.

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.

Bottenberg R.A., Ward J.H., 1963, Applied multiple linear regression. 6570th Personnel Research Laboratory, Aerospace Medical Division, Air Force

Bubbico R., Lee S., Moscati D., Paltrinieri N., 2020, Dynamic assessment of safety barriers preventing escalation in offshore Oil&Gas, Saf. Sci. 121, 319–330. DOI: 10.1016/J.SSCI.2019.09.011

Bucelli M., Paltrinieri N., Landucci G., 2017, Integrated risk assessment for oil and gas installations in sensitive areas, Ocean Eng. DOI: 10.1016/j.oceaneng.2017.12.035

Chen H., Moan T., Verhoeven H., 2008, Safety of dynamic positioning operations on mobile offshore drilling units, Reliab. Eng. Syst. Saf. 93, 1072–1090.

De Marchi B., Ravetz J.R., 1999, Risk management and governance:: a post-normal science approach, Futures 31, 743–757.

Duan X., 2018, Application of Artificial Intelligence in Evaluation and Management of SEC Oil and Gas Reserves, Chemical Engineering Transaction, 71, 925-930.

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. Tensorflow TM.

ISO, 2018. Risk Management. ISO 31000:2018.

Kaplan S., Garrick B.J., 1981, On The Quantitative Definition of Risk, Risk Anal. 1, 11–27.

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. DOI: 10.1016/j.coche.2016.07.006

Lee S., Landucci G., Reniers G., Paltrinieri N., 2019, Validation of Dynamic Risk Analysis Supporting Integrated Operations Across Systems, Sustain. DOI: 10.3390/su11236745

Matteini A., 2015, Human factors and dynamic risk analysis: a case-study in Oil&Gas drilling, Bologna, Italy.

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

Paltrinieri N., Dechy N., Salzano E., Wardman M., Cozzani V., 2012, 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. DOI: 10.1111/j.1539-6924.2011.01749.x

Paltrinieri N., Khan F., Amyotte P., Cozzani V., 2014, Dynamic approach to risk management: Application to the Hoeganaes metal dust accidents, Process Saf. Environ. Prot. 92. DOI: 10.1016/j.psep.2013.11.008

Paltrinieri N., Massaiu S., Matteini A., 2016, Human Reliability Analysis in the Petroleum Industry: Tutorial and Examples, Butterworth-Heinemann, pp. 181–192. DOI: 10.1016/B978-0-12-803765-2.00015-9

Patriarca R., Bergström J., Di Gravio G., Costantino F., 2018, Resilience engineering: Current status of the research and future challenges, Saf. Sci. 102, 79–100. DOI: 10.1016/j.ssci.2017.10.005

Patriarca R., Falegnami A., De Nicola A., Villani M.L., Paltrinieri N., 2019, Serious games for industrial safety: An approach for developing resilience early warning indicators, Saf. Sci. 118, 316–331. DOI: 10.1016/J.SSCI.2019.05.031

Samuel A.L., 1959, Some Studies in Machine Learning Using the Game of Checkers, IBM J. Res. Dev. 3, 210–229.

Sasaki Y., 2007, The Truth of the F-Measure, 2007.

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. DOI: 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, DOI: 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.