Reliability Engineering and System Safety 94 (2009) 855–860

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

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

Natural disaster risk analysis for critical infrastructure systems: An approach based on statistical learning theory ☆

Seth D. Guikema *

Department of Geography and Environmental Engineering, Johns Hopkins University, Baltimore, MD 21218, USA

ARTICLE INFO

Article history:
Received 25 June 2008
Received in revised form
12 September 2008
Accepted 17 September 2008
Available online 27 September 2008

Keywords:
PRA
Statistical learning theory
Infrastructure
Risk
Reliability
Data

ABSTRACT

Probabilistic risk analysis has historically been developed for situations in which measured data about the overall reliability of a system are limited and expert knowledge is the best source of information available. There continue to be a number of important problem areas characterized by a lack of hard data. However, in other important problem areas the emergence of information technology has transformed the situation from one characterized by little data to one characterized by data overabundance. Natural disaster risk assessments for events impacting large-scale, critical infrastructure systems such as electric power distribution systems, transportation systems, water supply systems, and natural gas supply systems are important examples of problems characterized by data overabundance. There are often substantial amounts of information collected and archived about the behavior of these systems over time. Yet it can be difficult to effectively utilize these large data sets for risk assessment. Using this information for estimating the probability or consequences of system failure requires a different approach and analysis paradigm than risk analysis for data-poor systems does. Statistical learning theory, a diverse set of methods designed to draw inferences from large, complex data sets, can provide a basis for risk analysis for data-rich systems. This paper provides an overview of statistical learning theory methods and discusses their potential for greater use in risk analysis.

1. Introduction

Probabilistic risk analysis (PRA) methods have historically been developed for situations in which there is little, if any, data about the operation of the system being analyzed. Two examples of typical problems addressed with PRA are estimating and managing the risk of accidents in nuclear power plants on the basis of fairly short operating histories (e.g., [1–3]) and estimating and managing the risk of failure of unique space missions (e.g., [3–5]). In situations like these, there is little, if any, hard data from the operation of the system being analyzed. Typically, the system is decomposed to the component level, and appropriate methods are used to estimate the reliability of the components based on a combination of component-level data and expert judgment. These component-level estimates are then rolled up to the system level through the use of appropriate techniques such as fault trees, event trees, and event sequence diagrams (e.g., [6,7]). This approach has worked well for many systems in a data-poor environment, and there are many extensive studies, books, and journal articles establishing and improving this approach. However, recent advances in information technology (IT) systems such as SCADA (Supervisory Control and Data Acquisition) and OMS (Outage Management Systems) for water and electric power systems, respectively, together with increased interest in estimating and managing the reliability of large-scale, networked infrastructure systems in areas prone to natural disasters make it necessary to examine other approaches as well.

The 1997 report from the President's Commission on Critical Infrastructure Protection focused attention on estimating and managing risk to critical infrastructure, those systems that are critical to life, livelihood, and governance in the US. While diverse, critical infrastructure systems share common characteristics. They typically cover extensive geographic scales, often spanning multiple geographic regions and, in some cases, multiple states. They are comprised of a very large number (e.g., tens of thousands or hundreds of thousands) of individual components often arranged as a network. Modern infrastructure systems are also typically monitored in some way through the use of IT systems. The geographic extent, large number of components in critical infrastructure systems, and increasingly large and detailed data sets about the performance of infrastructure systems during adverse events present difficulties for applying traditional PRA

☆ This paper is based on a paper presented at Probabilistic Safety Assessment and Management 9 in Hong Kong in May 2008.

* Tel.: +1 410 516 6042.
E-mail address: sguikema@jhu.edu

0951-8320/$ - see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ress.2008.09.003

methods at any level of detail to all but small, localized portions of larger systems.

While a number of approaches have been developed for using variations of PRA for analyzing infrastructure systems, these approaches are limited in their ability to probabilistically assess the reliability of critical infrastructure systems during natural disasters. In early work, Ezell et al. [8] advocated the use of a traditional PRA-based approach for infrastructure risk analysis based on using event trees and fault trees together with expert assessments to assess the risk of an attack against specific portions of a water supply system. This approach was developed in substantially more detail by Garrick et al. [9] and then built on to include an explicit uncertainty model by Torres et al. [10]. Paté-Cornell and Guikema [11] developed a PRA approach using a combination of Bayesian belief networks and game theory to model intelligent attacks against infrastructure and related targets. However, the approaches of [8–11] focus on modeling terrorist attacks against infrastructure, not natural disasters. They also do not directly utilize the wealth of information available about infrastructure system performance during past adverse events. Apostolakis and his colleagues [12–14] developed an approach for ranking the elements of infrastructure networks according to their expected disutility to a group of stakeholders. While this approach provides a valuable screening tool for ranking the potential importance of infrastructure elements, it does not provide a probabilistic estimate of the performance of infrastructure systems during natural disasters, and it does not directly utilize the wealth of information available about infrastructure system performance during past adverse events. It is also based on aggregation of the infrastructure networks, leading to a loss of fidelity [13] and it assesses the impacts of only single-element failures due to computational limitations for realistic systems [13,14]. Natural disasters are likely to lead to simultaneous failure of many system elements, requiring that a risk analysis approach for these events be capable of modeling these multi-element failures. In a similar line of research, Lambert and Sarda [15] developed an approach for identifying interactions among terrorist attacks scenarios and elements of infrastructure networks. While this approach provides useful input to a terrorist risk analysis for infrastructure systems, it is not directly applicable to natural disaster risk analysis, and it does not make use of the wealth of information available about infrastructure system performance during past adverse events.

The extensive and growing use of IT systems to capture data about infrastructure performance (e.g., traffic delays in a transportation system, loss of power in an electric power network, and signs of electronic intrusion in a banking system) has provided rich data sets on which risk analyses can be built for these systems. However, using these data effectively is difficult for a number of reasons. Data about infrastructure system performance generally come from a variety of past operating conditions that may not fully reflect situations of interest in the future, the relationship between the measures of the operating conditions (e.g., the loads placed on the different portions of the system) and system performance may be complicated and poorly understood, and the data sets themselves may be very large, making it difficult to draw clear insights from them. The problem in such a data-rich situation is one of learning—learning about system performance in the future based on a large data set about past performance. The field of statistical learning theory offers advantages in these situations that can help risk analysts and complement the existing set of infrastructure risk analysis tools. This paper introduces statistical learning theory, provides overviews of several examples of the application of tools from this field to the problem of risk analysis for infrastructure systems, and discusses the potential for further application and development of these approaches for risk analysis.

2. Statistical learning theory and its relationship to PRA

Statistical learning theory deals with learning from data [16–18]. This paper will deal in particular with supervised learning, a subset of statistical learning theory which is an approach for situations in which we have a record of outcome measures together with their associated explanatory measures, perhaps from past operation of the system. One example would be a situation in which we have an outcome measure we would like to predict (e.g., the number of customers without electric power in a given geographic area after the occurrence of a hurricane) and a number of other measures (e.g., wind speed, precipitation, and power demands) that may help us predict the measure of interest. Letting the matrix $\mathbf{X}$ denote the explanatory variables and y represent the outcome measure of interest, the goal of statistical learning theory is to estimate the relationship $y = f(\mathbf{X})$ where $f(\mathbf{X})$ is an unknown function of the explanatory variables. Note that $f(\mathbf{X})$ need not be a deterministic function (i.e., a function that does not incorporate any notion of uncertainty in y given $\mathbf{X}$). Indeed, in many situations of interest to risk analysts, $f(\mathbf{X})$ will contain considerable uncertainty. Statistical learning theory approaches this problem as one of using a set of training data (e.g., information about past system performance) to learn the form and parameters of a model approximating $f(\mathbf{X})$ in a way that will hopefully provide good predictions of future system performance. There is generally no direct use of information about the system's structure or the functional relationships between the components of the system as there is in PRA.

The problem of estimating $y = f(\mathbf{X})$ is also addressed when conducting a PRA. In conducting a PRA for a system, the outcome measure (y) is whatever risk measure is of interest, with common measures including the probability of successful system operation on the next demand cycle, the number of failures in a given number of trials, the time before the next failure, or the expected utility due to system use. If there is information, either hard data or expert opinion, about the operating conditions of the system, this can be used through appropriate methods such as fault trees and event trees to develop the conditional probability distributions for the outcome measure y. This is usually done by using a functional analysis of the system to decompose the system into its constituent functional units, analyzing the performance of each functional unit (i.e., estimating y_i, the performance of the ith functional unit), and then combining these unit-level estimates to obtain an estimate of the overall system performance. Thus, PRA assumes that the form of the relationship $y = f(\mathbf{X})$ is determined through a functional analysis of the system.

The key difference between statistical learning theory and PRA is the assumption made about the form of the relationship $y = f(\mathbf{X})$. Statistical learning theory assumes that the form of this relationship is to be estimated from the data, whereas PRA assumes that the form of this relationship is to be estimated based on knowledge of the logic of the system's performance. This may take the form of a simulation model for estimating the system's performance as in [14]. This difference is critical in risk analysis for large-scale infrastructure networks such as multi-state electric power distribution systems. The size and complexity of these systems make the development of full PRA models addressing common-mode failures due to natural hazards based on the decomposition approach prohibitively difficult. While it may be possible to partially decompose the system (e.g., decompose a power system down to the level of individual substations rather than down to the level of individual line segments) and then assess reliability at this less detailed level of decomposition, this can still be prohibitively difficult for large systems while also not providing detailed information about the link between system performance and the characteristics of individual system components. Furthermore, for realistic systems of interest, the large number of possible combinations of multi-component failures generally makes traditional PRA approaches relying on assessing the impacts of enumerated failure scenarios (e.g., [8–14]) computationally infeasible. Statistical learning theory provides an alternate and complimentary approach. It does not replace PRA-based methods, which still provide valuable insight into component importance. Rather, statistical learning theory can provide a system-level assessment of the impacts of natural disasters on infrastructure reliability.

3. Statistical learning theory methods

Statistical learning theory is a broad field, covering many different methods and approaches. Perhaps the simplest and most well-known approach that can be classified as a supervised learning method is linear regression. Linear regression assumes that the relationship between y and the elements of X, x_j, is linear. That is, linear regression assumes that the form of the model is $y_i = \sum_j^n \beta_j x_{ij}$, and, after making some additional assumptions about the form of the errors (e.g., that the they follow a zero-mean Gaussian distribution), and the fitting criteria (e.g., minimizing the sum of squared errors), the past data are used to find the values of the β parameters that yield the best fit based on the past data. Most risk analysts and reliability engineers are familiar with linear regression. However it is not a particularly flexible approach, and it does not provide particular good predictions of system performance in many situations. Fortunately, statistical learning theory provides a much richer, more flexible set of models for supervised learning. In the following sections a selection of these models are introduced together with examples of their use in problems of direct relevance to risk and reliability analysts. The purpose here is not to provide an exhaustive overview of the methods of statistical learning theory. Rather, the intent of the following discussion is to demonstrate the breadth of modeling approaches and their use for modeling infrastructure system performance.

3.1. Parametric regression models

Parametric regression models provide a class of approaches in which the function $y = f(\mathbf{X})$ is approximated by assuming a particular parametric form for the relationship. The linear regression model discussed above is one simple approach. Nonlinearity can be accommodated in the linear regression model by including transformations of some or all of the elements of $\mathbf{X}$ as explanatory variables. Discrete responses (i.e., discrete-valued y) can be accommodated through the use of a generalized linear model (GLM), a useful class of models for risk and reliability analysis where the events of interest are often either binary (e.g., failure/success) or counts (e.g., number of failures or successes).

A GLM consists of three components. The first component, the random component, specifies the behavior of the response variable for a fixed set of predicting variables by allowing the use of any distribution from the exponential family as [19]

$$f_Y(y;\theta;\phi) = \exp\{(y\theta - b(\theta))/a(\phi) + c(y,\phi)\} \tag{1}$$

where θ is the natural parameter and ϕ is the scale parameter. The second component, the systematic component, of a GLM specifies the predicting variables and the relationship between them and the link parameter η is

$$\eta = \beta_0 + \sum_i \beta_i x_i. \tag{2}$$

The final component links the systematic and random components. For example, the mean μ may be related to the natural parameter θ by $\mu = b'(\theta)$ [19,20]. This link component generally is of the form:

$$g(\mu) = \eta \tag{3}$$

where g is referred to as the link function. A simple example of a GLM is the Poisson GLM given by Eqs. (4) and (5):

$$y \sim \text{Poisson}(\lambda) \tag{4}$$

$$\log \lambda = \beta_0 + \sum_i \beta_i x_i. \tag{5}$$

Examples of the use of GLMs for risk analysis include modeling power outages during hurricanes [21,22].

The goal of the models of Liu et al. [21] and Han et al. [22] is to probabilistically estimate the number and geographic distribution of power outages during a hurricane 1–3 days before the hurricane makes landfall. This estimate of the reliability of the power system during an approaching hurricane would then provide a basis for pre-storm planning for post-storm recovery efforts. These models utilized the GLM framework. The service area of the utility company was decomposed to a number of smaller geographic areas. Data about power outages during past hurricanes together with the values of predictor variables were collected at the level of the individual geographic areas, and negative binomial GLMs were fit to these past data. A number of explanatory variables (the elements of $\mathbf{X}$) were used in the models, including measures of the wind speed experienced in each geographic area during the past hurricanes, the number of poles, transformers, switches, and other system components in each geographic area, land use measures, and a number of other variables related to hurricane and geographic characteristics. Han et al. [22] built from Liu et al. [21] to include a number of measurable hurricane and climatic characteristics (e.g., rainfall, soil moisture, hurricane wind radius) that help to increase the predictive accuracy of the GLMs. The output from models such as these is an estimate of the number of power outages that can be expected in each geographic area of a utility company's service area during an approaching hurricane.

Parametric regression models are relatively easy to use, they have a solid theoretical basis, and they can provide useful estimates in some situations. It is not surprising then that they are widely used in risk and reliability analysis. Examples include modeling traffic accident risk [23,24], modeling the reliability of water distribution systems [25], and modeling electric power system reliability [26,27]. While these models can be useful in many situations, the parametric assumptions, particularly the common assumption of additive linearity in the relationship between the response variable and the explanatory variables, can be restrictive. A number of semi-parametric and non-parametric approaches have been developed to relax the assumptions inherent in parametric models.

3.2. Semi-parametric and non-parametric regression models

There are a wide variety of semi-parametric and non-parametric regression models, and providing a summary of all of these approaches is beyond the scope of this paper. The goal here is to introduce the idea behind these methods by examining two, generalized additive models (GAMs) and multi-variate adaptive regression splines (MARS). It is hoped that this limited overview will give a sense of the flexibility and diversity of these approaches.

One extension of the GLM framework is to relax the assumption of linearity in the link function shown in Eq. (2), while keeping the remainder of the GLM framework intact. Relaxing this assumption through the use of smoothing splines over the individual variables or combinations of variables specified prior to fitting the model leads to a GAM. The link function becomes

$$\eta = \beta_0 + \sum_i f_i(x_i) \tag{6}$$

where $f_i(x_i)$ is a smooth over the ith element of $\mathbf{X}$ (see, for example, [20]). Cubic regression splines and thin-plate splines are common smoothing functions used in GAMs. The smoothes are adaptive in the sense that the shape of the function is not specified prior to the analysis. Rather, the algorithm used to fit a GAM chooses the shape of the function to best fit the past data. GAMs are more difficult to fit than GLMs due to the need to fit a set of smoothing functions rather than a set of simple regression parameters. In addition, while the increased flexibility of GAMs does mean that they can fit a given data set better than GLMs in many cases, it also increases the risk of 'overfitting' the data—fitting the model to features in the historic data that may not be present in future situations. This can lead to poor predictive accuracy from GAM models for situations outside of the fitting data set, making thorough model validation essential when GAMs are to be used. GAMs have been used in risk and reliability analysis, though not as widely as GLMs. For example, Han et al. [28] developed a GAM for estimating power outages during hurricanes using the same data set used in [22]. In this modeling effort, modeling the non-linearity in the data set did significantly improve the predictive accuracy of the model relative to GLMs, as measured through hold-out validation testing. This is not necessarily always the case.

Another class of models of interest is MARS models [29]. A MARS model is a non-parametric approach that allows for complex interactions between the explanatory variables in a way that allows the fitting algorithm to determine which explanatory variables should be modeled as interacting based only on the data. That is, the analyst does not pre-specify the interactions, unlike in the GAM framework. A MARS model consists of a family of interacting basis functions $B(x, \theta)$ where x is an element of $\mathbf{X}$ and θ is a 'knot point', a point in the domain of $\mathbf{X}$ that determines a change in behavior of the basis function. The types of functions that can be used as basis functions is diverse, with some examples being piece-wise linear functions, cubic or tensor-product splines, and wavelets. The full MARS model is given by

$$f_i(x_i) = \sum_j \beta_j B(x_j, \theta) \tag{7}$$

where β_j is a regression parameter for the jth basis function. Details on the process of fitting these models are available in [29]. The high degree of flexibility available in MARS models allows them to fit historic data very well. However, this increases the chances of overfitting if careful model validation is not conducted.

MARS models have not been widely used in risk and reliability analysis. One related example is the use of MARS for modeling traffic accidents [30]. This work used a MARS model to estimate the number of traffic accidents in different geographic areas in the state of Texas. Similar methods could be developed for estimating the reliability measures for complex, distributed infrastructure systems during natural disasters.

3.3. Machine-learning, artificial intelligence, and pattern-matching methods

The methods discussed so far have all been regression approaches or generalizations of regression approaches. However, a large class of methods exist which are not directly related to or generalizations of regression-based approaches. These methods have been developed in a variety of fields, and they go by a variety of names. Here they will be referred to generically as machine-learning, artificial intelligence, and pattern-matching methods.

One of the more widely known and used artificial intelligence approaches is artificial neural networks (e.g., [31]). These models seek to approximate the relationship $y = f(\mathbf{X})$ by mimicking the functioning of neurons in the brain. The model consists of layers of interconnected nodes with weights on each of the links. These weights are iteratively adjusted to match the output of the network (estimated values of y) with the measured values of y based on the input values in $\mathbf{X}$. The goal of artificial neural networks then is to match the pattern in the training data based on a flexible, non-parametric model. The trained network is then often used for predicting future system states. Artificial neural networks have been used in risk and reliability analysis in a number of areas such as nuclear power plant reliability [32], software fault detection [33], and more general reliability and availability modeling [34].

A second example of a machine-learning approach is support vector machines (SVMs). SVMs are a member of the larger family of linear classifiers. The goal of linear classification methods is to classify a set of data, each element of which can belong to one of the two sets. Each element of the data has an associated set of variables (the elements in $\mathbf{X}$) that can help with the classification. SVMs take this one step farther and seek to also find maximum separation between the classes with the classifier. There are many possible applications of SVMs in risk and reliability analysis. A few examples are classifying events (e.g., transients in nuclear power plants) as safety problems or not based on measurements of observable characteristics of these events, classifying a potentially contaminated site as safe or not based on sparse observations of sub-surface conditions, and classifying the performance of a system as acceptable or not based on easily measured system characteristics. Despite the apparent usefulness of SVMs for risk and reliability analysis, SVMs are not widely used in risk and reliability analysis literature. Two examples where they have been used is for classifying transients in nuclear power plants [35] and for approximating computationally expensive performance functions when evaluating the reliability of a complex system [36]. SVMs are a promising approach for classification problems in infrastructure risk and reliability analysis.

There are many other machine-learning, artificial intelligence, and pattern-matching approaches available, and a full review of these methods is beyond the scope of this paper. However, the overview of the two approaches given above highlights the key characteristics of these methods. They are highly flexible, non-parametric, and seek to approximate $y = f(\mathbf{X})$ based only on the patterns observed in the past data. This flexibility makes them advantageous for modeling complex infrastructure problems. At the same time, overfitting is a potential problem with these models, and careful model validation and testing are needed.

3.4. Selection of explanatory variables

A critical step in using statistical learning theory methods for risk analysis is the selection of the explanatory variables to use as input to the model. These variables form the basis for developing the predictive model, and selecting the proper explanatory variables depends on a sound knowledge of the problem. The standard approach is to use as input any variable that is both thought to potentially be relevant for predicting system performance and is measurable for future applications of the model. The model fitting process for many, though not all, statistical learning theory methods then provide information that can be used to reduce the set of input variables so that only those that help to improve the fit of the model to past data are included. Having access to the proper expertise for the problem being analyzed is important to this creative process. Without access to experts who understand the problem from different perspectives, important variables are likely to be overlooked. An example from the author's experience serves to illustrate this point.

Selecting explanatory variables to use in developing the GLM-based hurricane power outage prediction models of Han et al. [22] was a challenging, iterative process that depended on having access to the proper experts. Earlier work by Liu et al. [21] in developing GLM-based hurricane power outage prediction models, the first rigorous regression modeling done for this problem, had utilized a set of explanatory variables that included information about the power system such as the number of transformers and the miles of overhead line in an area, basic geographic information such as land use classifications, and information about wind speeds during past hurricanes. However, there were binary indicator variables for past hurricanes (e.g., "1" for hurricane A and "0" for the other hurricanes in the data set), variables that were not measurable prior to a future hurricane. This limited the applicability of the models to future hurricanes. The two main goals in developing the models of Han et al. [22] were to improve the predictive accuracy of these types of models and to use only explanatory variables that could be measured or easily estimated in advance of an approaching hurricane. To develop these improved models the authors collaborated with a climatologist and geographer to better understand hurricanes and the measurable characteristics of hurricanes and local geography that might be useful as explanatory variables. This led to the inclusion of additional climate variables that included soil moisture levels at three different depths for different time frames before storms and long-term precipitation patterns among other factors. The authors also worked with structural engineers to better understand the factors that affect the stability of poles during hurricanes, and this led to a particular emphasis on the soil moisture variables as potential determinants of pole foundation stability during high winds. Finally, the authors worked with engineers from the utility providing the historic hurricane outage data to develop a more complete set of information about the power system, including additional information about poles, transformers, switches, and related information. While some of these variables were shown to not yield significant improvements in model fit, others did. Overall, the models of Han et al. [22] offered an improved fit to the data and used only measurable variables. In ongoing work the authors are refining the set of variables used as input for the models. This example serves to illustrate that selecting the proper explanatory variables to use is (1) critical to the fit and predictive accuracy of the models, (2) an iterative, creative process, and (3) dependent on working with the right experts. As with PRA, expert knowledge plays a key role. However, this role is different. With statistical learning theory methods expert knowledge focuses mainly on selecting the right input and then selecting and implementing appropriate statistical learning theory methods.

4. Discussion and summary

Large-scale infrastructure systems are complex and geographically distributed, making the development of traditional PRA models based on a functional decomposition of the system problematic. There is also a wealth of information about past performance available for many infrastructure systems. While there will continue to be important classes of problems for which PRA is the best available tool, the growing use of IT systems to monitor infrastructure networks and record data pertaining to their performance during adverse events makes it important to explore and develop data-driven methods that can compliment PRA for estimating infrastructure system performance and reliability on the basis of large, complex data sets.

Statistical learning theory consists of a set of methods appropriate for this problem. These methods explicitly make the best available use of the wealth of performance data available for many infrastructure systems. They also model the system-wide impacts of natural disasters rather than only single-element failures, provided that the past data include events with these characteristics. Statistical learning theory methods, when applied carefully and appropriately, can yield a truly systems-level assessment of the reliability of infrastructure systems during natural disasters. As with PRA, this provides important support for risk management decision-making both immediately before a forecast event such as a hurricane and over the longer term. However, it is critical to point out that the conclusions from these methods are no better than the data input to them. If the data used to develop the models are not representative of the potential future disasters being analyzed, then the forecasts of the method should be treated with caution. For example, if the past performance data include performance data from only weak hurricanes (e.g., only category 1 storms) yet the impact of a strong hurricane on system performance (e.g., a category 5 storm) is being assessed, the model forecasts of performance must be treated cautiously. Attempting to use statistical learning theory models for predictions for conditions outside of the conditions present in the training data is one of the main errors that must be avoided with these types of methods. Similarly, care must be taken to thoroughly validate the models prior to using them for estimating performance. It is easy to over-fit some types of statistical learning theory models, particularly the more flexible non-parametric regression, machine-learning, artificial intelligence, and pattern-matching methods. An over-fit model may fit the past data very well encouraging naive confidence in the model if it is not thoroughly validated. This same model may lead to inaccurate predictions of future performance. If statistical learning theory methods are to be used to complement PRA-based methods, careful validation studies are needed for the resulting models. Similarly, methods must be selected such that the assumptions underlying the method such as linearity (or not) and additivity (or not) in the response are appropriate. Guikema and Coffelt [37] provide a discussion of these assumptions. Still, provided that the data available are appropriate, proper methods are selected, and the models are developed carefully, statistical learning theory methods can provide a strong complement to PRA-based approaches for natural hazard risk analysis for infrastructure systems. They could also potentially be used for real-time monitoring of infrastructure systems to detect abnormal system performance. Several research groups have ongoing work in this area.

This paper has provided an introduction to statistical learning theory methods and their use in risk and reliability analysis for infrastructure systems. These methods can provide a strong complement to PRA-based approaches. While some of these methods are beginning to be used in some areas of risk and reliability analysis, further development and refinement of these methods will help to advance risk and reliability analysis for infrastructure systems.

Acknowledgements

This work was partially funded by a grant from the National Science Foundation (EECS-0725823). This support is gratefully acknowledged.

References

[1] Rasmussen N. An assessment of accident risks in US Commercial Nuclear Power Plants. Nuclear Regulatory Commission Report. s.l.: Nuclear Regulatory Commission, WASH-1400; 1975.

[2] Lewis HW, Budnitz RJ, Rowe WD, Kouts HJC, von Hippel F, Loewenstien WB, et al. Risk assessment review group report to the US Nuclear Regulatory Commission. IEEE Trans Nucl Sci 1979;NS-26:4686–90.

[3] Apostolakis GE. How useful is quantitative risk assessment? Risk Anal 2004;24:515–20.

[4] Dillon RL, Pate-Cornell ME, Guikema SD. Programmatic risk analysis for critical engineering systems under tight resource constraints. Oper Res 2003;51:354–70.

[5] Pate-Cornell ME, Dillon R. Probabilistic risk analysis for the NASA space shuttle: a brief history and current work. Reliab Eng Syst Saf 2001;74: 345–52.

[6] Modarres M. Risk analysis in engineering: techniques, tools, and trends. Boca Raton, FL: CRC Press; 2006.

[7] Pate-Cornell ME. Fault tress vs. event trees in reliability analysis. Risk Anal 1984;4:177–86.

[8] Ezell BC, Farr JV, Wiese I. Infrastructure risk analysis model. J Infrastructure Syst 2000;6(3):114–7.

[9] Garrick BJ, Hallo JE, Kilger M, McDonald JC, O'Toole T, Probst PS, et al. Confronting the risk of terrorism: making the right decisions. Reliab Eng Syst Saf 2004;8:129–76.

[10] Torres J, Brumbelow K, Guikema SD. Risk classification and uncertainty propagation for virtual water distribution systems. Reliab Eng Syst Saf, under review.

[11] Paté-Cornell ME, Guikema SD. Probabilistic modeling or terrorist threats: a systems analysis approach to setting priorities among countermeasures. Mil Oper Res 2002;7(4):5–23.

[12] Apostolakis GE, Lemon DM. A screening methodology for the identification and ranking of infrastructure vulnerabilities due to terrorism. Risk Anal 2005;25(2):361–76.

[13] Michaud D, Apostolakis GE. Methodology for ranking elements of water-supply networks. J Infrastructure Syst 2006;12(4):230–42.

[14] Koonce AM, Apostolakis GE, Cook BK. Bulk power risk analysis: ranking infrastructure elements according to their risk significance. Electr Power Energy Syst 2008;30:169–83.

[15] Lamber JH, Sarda P. Terrorism scenario identification by superposition of infrastructure networks. J Infrastructure Syst 2005;11(4):211–20.

[16] Hasti T, Tibshirani R, Friedman J. The elements of statistical learning theory. New York: Springer; 2001.

[17] Vapnik VN. An overview of statistical learning theory. IEEE Trans Neural Networks 1999;10:988–99.

[18] Xuegong Z. Introduction to statistical learning theory and support vector machines. Acta Automatica Sin 2000;26:32–42.

[19] Cameron AC, Trivedi PK. Regression analysis of count data. Cambridge: Econometric Society Monographs, Cambridge University Press; 1998.

[20] Hastie T, Tibshirani R. Generlized additive models. London, Boca Raton, FL: Chapman & Hall, CRC Press; 1990.

[21] Liu H, Davidson RA, Rosowsky DV, Stedinger JR. Negative binomial regression of electric power outages in hurricanes. J Infrastructure Syst 2005;11:258–67.

[22] Han S, Guikema SD, Quiring SM, Lee K, Rosowsky D, Davidson RA. Estimating the spatial distribution of power otuages during hurricanes in the Gulf Coast region. Reliab Eng Syst Saf 2008 [accepted February 2008].

[23] Lord D, Washington SP, Ivan JN. Poisson, Poisson-gamma, and zero-inflated regression models of moter vehicle crashes: balancing statistical fit and theory. Accid Anal Prev 2005;37:35–46.

[24] Lord D, Guikema SD, Geedipally S. Application of the Conway–Maxwell–-Poisson generalized linear model for analyzing motor vehicle crashes. Accid Anal Prev 2008;30(3):1123–34.

[25] Yamijala S, Guikema SD, Brumbelow K. Statistical estimation of water distribution system pipe reliability. Reliab Eng Syst Saf 2009;94(2):282–93.

[26] Radmer DT, Kuntz PA, Christie RD, Venkata SS, Fletcher RH. Predicting vegetation-related failure rates for overhead distribution feeders. IEEE Trans Power Delivery 2002;17:1170–5.

[27] Guikema SD, Davidson RA, Liu H. Statistical models of the effects of tree trimming on power system outages. IEEE Trans Power Delivery 2006;21:1549–57.

[28] Han S, Guikema SD, Quiring SM. Improving the accuracy of hurricane power outage predictions using generalized additive models. Risk Analysis [under review for publication].

[29] Friedman J. Multivariate adaptive regression splines. Ann Stat 1991;19:1–141.

[30] Song JJ, Ghosh M, Miaou S, Mallick B. Bayesian multivariate spatial models for roadway traffic crash mapping. J Multivar Anal 2006;97:243–73.

[31] Haykin S. Neural networks: a comprehensive foundation. Upper Saddle River: Prentice Hall; 1994.

[32] Santosh TV, Vinod G, Saraf RK, Ghosh AK, Kushwaha HS. Application of artificial neural networks to nuclear power plant transient diagnosis. Reliab Eng Syst Saf 2007;92:1468–72.

[33] Hu QP, Xie M, Ng SH, Levitin G. Robust recurrent neural network modeling for software fault detection and correction prediction. Reliab Eng Syst Saf 2007;92:332–40.

[34] Rajpal PS, Shishodia KS, Sekhon GS. An artificial neural network for modeling reliability, availability and maintainability of a repairable system. Reliab Eng Syst Saf 2006;91:809–19.

[35] Rocco CM, Zio E. A support vector machine integrated system for the classification of operation anomalies in nuclear components and systems. Reliab Eng Syst Saf 2007;92:593–600.

[36] Rocco CM, Moreno JA. Fast Monte Carlo reliability evaluation using support vector machine. Reliab Eng Syst Saf 2007;76:237–43.

[37] Guikema SD, Coffelt JP. Practical considerations in statistical modeling of count data for infrastructure systems. J Infrastructure Syst 2008, in press.