Artificial Intelligence for Natural Hazards Risk Analysis: Potential, Challenges, and Research Needs

Seth Guikema, University of Michigan, sguikema@umich.edu

Abstract

Artificial intelligence (AI) methods have seen increasingly widespread use in everything from consumer products and driverless cars to fraud detection and weather forecasting. The use of AI has transformed many of these application domains. There are ongoing efforts at leveraging AI for disaster risk analysis. This paper takes a critical look at the use of AI for disaster risk analysis. What is the potential? How is the use of AI in this field different from its use in non-disaster fields? What challenges need to be overcome for this potential to be realized? And what are the potential pitfalls of an AI-based approach for disaster risk analysis that we as a society must be cautious of?

1. Introduction

Natural hazards pose significant risks throughout the world. They are among the deadliest disasters. These events cause significant economic damage as well, with losses from a large tropical cyclone impacting a developed nation approaching or, at times, exceeding $100 billion U.S. dollars.

Risk analysis is, in broad terms, a systematic process aimed at understanding the nature of risk in a given situation and expressing the risk together with the underlying knowledge base (SRA, 2015). It is usually thought of as being comprised of at least risk assessment, risk management, risk communication, and risk governance among other aspects. Risk assessment, the process of understanding and characterizing the risk, often mathematically, is the primary focus of this paper.

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1111/risa.13476.

Traditionally, risk assessment is thought of as answering the following three questions: (1) what can go wrong? (2) with what likelihood, and (3) with what consequences? (Kaplan and Garrick, 1981). More recently, there has been an increasing focus on the background state of knowledge underlying these assessments and the uncertainty in these assessments (e.g., Aven 2013, 2017).

Risk analysis is critical for natural hazards, and a number of different risk analysis methods exist for assessing risk to communities from natural hazards. One set of approaches are fragility-based models such as the widely-used HAZUS model that are based on simulating physical loads to systems, estimating asset-level damage through fragility curves, and then estimating system performance, losses, and deaths (e.g., Kircher et al. 2006, Schneider and Schauer 2009, Winkler et al., 2010). A second set of approaches are based on machine learning and artificial intelligence methods (e.g., Guikema et al. 2014, Quiring et al. 2014, Baroud and Barker 2018, Shashaani et al. 2018). These methods will be discussed more below, but briefly, they use past data and hazard loading information, information about the system and antecedent conditions, and other information to train a machine learning model that is then used to estimate system performance or losses for future events. A third type of approach that has seen widespread use in the academic literature, but less in practice, is network theory based methods (e.g., Dueñas-Osorio and Vemuru 2009, Wang et al. 2013). These approaches focus on leveraging the topology of the system to estimate performance under hazard loading in a more computationally efficient manner, though the accuracy of the results has been questioned (LaRocca et al. 2015).

The primary focus of this paper is on artificial intelligence, machine learning, and statistical methods, referred to as AI methods in the remainder of this paper. These methods are becoming increasingly used in practice, and significant research advances have been made. But a critical, foundational look at these methods and their use for natural hazard risk assessment is needed. They can be very valuable and useful, but they are not the panacea some present them to be.

2. Artificial Intelligence Methods and Their Expanding Reach

Artificial intelligence, machine learning, and statistical learning theory are closely related bodies of knowledge and methods that can be separated into two separate classes of methods, supervised learning and unsupervised learning. Supervised learning seeks to estimate an unknown relationship f( ) between a set of explanatory factors, **x**, and one or more response variables, y. That is, supervised learning seeks to estimate the relationship y = f(**x**). This relationship is then often used to make predictions of the value(s) or probability distributions of y' given a new set of inputs **x'** in some unobserved, likely future, situation (e.g., predicting power outages prior to hurricane landfall as in Guikema et al. 2014). The learned relationship can also be used to draw insights into the influence of different explanatory factors in **x** on y (e.g., Rivero-Calle et al. 2015). It should be noted that supervised learning methods can produce either point estimates or probability distributions as predictions, with approaches such as quantile regression forests (e.g., Kabir et a. 2019) and Bayesian belief networks trained with past data (e.g., Francis et al. 2014) being examples of probabilistic methods. Unlike supervised learning, unsupervised learning does not involve y. Instead, it seeks to understand and model the relationships between the elements of **x** to provide insight into the problem or to help support further analysis.

AI and related methods have proliferated in practice across many fields, including, among others, credit card fraud detection (e.g., Chan et al. 1998, Maes et al. 2002), automated vehicle control (e.g., Pomerleau and Jochem 1996, Burton et al. 2017), and climate science (Badr et al. 2014, Rivero-Calle et al. 2015, Tripahi et al. 2006). The use of these methods in practice has also increased dramatically with many companies, from start-ups to long-established companies branding or rebranding themselves as data science or AI companies. These methods have found significant success in prediction in many of these fields.

Why have AI and related methods become so widely-used? Underlying their explosive growth have been a rapid proliferation of data, the development of efficient computational hardware that can now handle large data volumes together with complicated models, and demonstrated success of these methods in application domains for which large financial outcomes are at stake such as credit card fraud detection. Many of the most successful application domains can be characterized as situations in which there is a large amount of data from repeated occurrences of the same or at least highly similar situations. Consider the challenge of recognizing stop signs algorithmically based on sensors in an automated vehicle, an area where AI models have had success (e.g., Huang et al. 2016). It is relatively easy to construct a large training set of stop signs as viewed by the vehicle's sensors under different weather and lighting conditions. Stop signs do not change over time, so if a sufficiently diverse training set can be created, one would reasonably expect AI methods to work well for this problem. Similarly, credit card fraud detection has been a successful application of AI and statistical models because (1) large training sets of transactions, labeled as fraudulent or not, can be created and (2) these data are representative of future data streams that will be evaluated. These two characteristics, a large corpus of diverse training data and future conditions that are well-represented within the training data, are critical to the success of AI methods (Guikema, 2009). If these two criteria are not met, the trained AI model(s) are being asked to do something they have not been trained to do - make predictions for situation for which they have no representative data on which to make a prediction.

3. AI for Natural Hazards Risk Analysis

AI methods have been used in both research and practice for natural hazards risk assessment. This existing work has generally focused on estimating one or more of the following: (1) the physical loading due to the hazard given occurrence of the hazard, or (2) physical damage or loss of system functionality given hazard loading. Examples of each are provided below. The intention is not to be exhaustive but to give representative examples.

One example of the use of AI for estimating properties of the physical hazard itself is in the area of flood modeling. Several start-ups are now working on approaches that ingest weather forecasts and predict the timing and extent of riverine flooding on fine spatial scales across watershed-scale domains using a combination of physical models and validated AI methods. A similar example is recent research that aims to estimate flooding without the use of computationally expensive fluid dynamics models (e.g., Liong and Sivapragasam 2002, Mosavi et al. 2018). Often physical models are used to developed a training data set and then AI methods are trained and validated to predict what the flood map would look like given information about the storm and the geographic area. Models such as these that aim to predict hazard loading are, by themselves, not a full risk assessment. However, they can provide valuable information for a risk assessment by estimating physical hazards in a more computationally efficient manner.

The area where AI methods have had their largest use in natural hazards risk assessment is in estimating either damage or loss of system function given hazard loading. For example, my collaborators and I have developed AI-based methods for estimating the spatial distribution and total number of power outages due to adverse weather events based on weather forecasts (e.g., Han et al. 2009, Guikema et al. 2014). AI-based models have similarly been developed to estimate building damage given the occurrence of an earthquake (e.g., Suryanita and Adnan 2012). These types of models take as input a spatial field of hazard loading (e.g., maps of predicted wind speeds, soil moisture levels, and other loading measures measures for a hurricane or a map of a ground motion measure for an earthquake) together with a set of information about the system, the area, and pre-event conditions. The model then estimates a spatial field of impacts (e.g., a map of power outages due to the approaching hurricane or a map of building damage states due to an earthquake). Again, these models are not, by themselves, a full risk assessment, but, like the hazard loading predictions, provide critical information for a full risk assessment.

In both of these types of applications of AI methods for natural hazards, the goal is to improve predictive accuracy and/or reduce computational burden relative to more traditional physics-based and engineering-based models. These models have seen use in practice recently. For example, a large portion of electric power utilities report the use of some type of power outage prediction model, and some states have required power utilities in their state to have an outage prediction model. Start-ups such as One Concern, Inc. have built businesses on providing predictions of hazard events. Clearly emergency managers, utility managers, and policy makers see value in AI-based models for natural hazards.

4. Characteristics of Settings in which AI Methods Have Been Successful

Across many domains, there are three key characteristics that have been critical to the success of AI methods. These are:

1. Large training sets exist with clearly labeled classes (for classification) or clearly measurable outcomes (for regression).
2. The training data used to train the model is representative of the future situations in which the model will be used to make predictions.

3. The relationships between the feature space (explanatory variables) and the response variable are the same in the future situations as in the training set.

Each of these will be discussed briefly below.

AI methods require large data sets to learn relationships between variables. This is true for models as diverse as single regression trees, ensembles such as random forests, and deep learning methods, though the data size considered large enough depends on the type of model with more complex models generally requiring more training data. These types of models are highly flexible, which is part of their appeal. However, this flexibility means that there are many different parameters of the model which must be fit to a given data set. If the available data is insufficient for the model structure used, insufficient learning will occur, leading to poor predictive accuracy. For example, AI methods have proven successful in detecting potentially fraudulent credit card transactions in part because of the extensive data sets available to credit card companies to use in training models. There is an important caveat needed here. Large data sets and models with complicated structures and many free parameters can lead to overfitting. In these cases, the many degrees of freedom of the model allow it to match variations in input and noise in relationships, patterns that may not be present in future situations. This then leads to reduced accuracy in future settings. Careful model validation and regularization is needed to reduce these problems. This is an especially important problem when the data is unbalanced, that is, when there are a far more zero (non-event) records than can be represented by standard distributions. I will return to these points in the discussion of validation testing of models.

Even if models are trained and validated properly, extensive data alone is not sufficient for an AI model to accurately predict outcomes in a future situation. The training data must also be representative of the future situation for which predictions are to be made. There are many dimensions to this. First, the future situation must be within the domain of the training data. For example, a model predicting flooding from hurricanes trained on data only from weak hurricanes cannot be expected to offer accurate predictions for a strong Category 5 storm. Similar, a building damage model trained on damage date from only weak and moderate earthquakes should not be trusted to yield good predictive accuracy for strong earthquakes. In applications such as image recognition and credit card fraud, this criteria is relatively easy to meet, provided a diverse enough set of training instances is used. This, however, may pose significant challenges for nature hazards. Strong events are thankfully rare for many types of hazards, limiting the training data available for the types of events we typically are most concerned with.

Even extensive training data representative of future conditions may be insufficient if the relationship between the explanatory features and the response variable is not static. That is, if the relationship between **x** and y that is learned by the model is not the relationship that will exist in the future, the model should not be expected to offer accurate predictions. A critical clarification is needed here, especially in the context of natural hazards. What needs to be stationary is f(**x**), that is, the relationship between **x** and y. The phenomenon, represented by **x**, does not necessarily need to be static, as long as f(**x**) is. For example, consider the problem of predicting power outages due to hurricane. The characteristics of the hurricane that impact the power system are described by **x**. An AI model learns the relationship, f(**x**) between **x** and y. As long as the first two conditions are met and the model achieves sufficient accuracy, it should work in future situations as long as f(**x**) has not changed, *even if the hazard itself is not stationary*. The critical point here is that f(**x**) must be stationary, not **x** itself. The hazard, **x**, may not be stationary (e.g., due to climate change) even while f(**x**) is stationary. How would f(**x**) not be stationary? If the system changed significantly, f(**x**) would not be stationary. For example, if an electric power utility buried many of their power lines or significantly strengthened their utility poles, a model trained with pre-change data would not offer accurate predictions for post-change data because f(**x**) changed. On the other hand, if the system had not changed substantially, f(**x**) may still be stationary, even if climate change is leading to changes in the frequency or intensity of storms; what is needed is that the response of the system to a given storm is stationary, not that storms themselves are stationary. As a corollary to this, climate change may differentially affect different aspects of **x**. For example, climate change may increase soil moisture levels (one set of variables in **x** in our prediction models) more than wind speeds. This is still fine for the use of AI methods as long the changes in **x** stay within the range of the training data and as long as there is sufficient training data available to represent the range of **x** experience in the future in the training of the model.

All sustained, successful uses of AI, in the sense that the model repeatedly makes accurate predictions, have met these three criteria. Lucky outcomes with one-off accurate predictions certainly occur. But for a model to have sustained accuracy, the conditions above are necessary.

5. Challenges in Using AI Methods for Natural Hazards Risk Analysis

The three characteristics above (sufficient data, data representative of future scenarios, and a stationary response relationship) pose significant challenges for the use of AI methods for natural hazards risk analysis.

The first and most obvious challenge is obtaining a sufficiently large and representative set of training data. For very rare hazards (e.g., strong solar storms that impact the power system) it may not be possible to assemble a data set that is large enough and contains a diverse enough set of events to support the use of AI methods. A variant of this problem that is more subtle is the case where there is a large set of training data, but the data set lacks diversity in terms of hazard characteristics. Consider a case in which one seeks to develop an AI-based approach for estimating earthquake damage to drinking water systems in the U.S. There are quite a few weak earthquakes for which one could conceivably gather water system damage and asset data. However, there are very few strong earthquakes in recent history from which one might be able to obtain data. The data set would be skewed towards weak events and only a small subset of the possible strong events would be covered. An AI-based approach would likely not be able to consistently offer strong predictive accuracy for strong events unless they were similar to the few in the training data set in terms of ground motions, depths, and other similar characteristics.

A second key challenge is the issue of validation, and this issue is related to the first. The traditional way to validate AI models is through holdout validation testing. A portion of the data is set aside, the model is trained on the remaining data, and then the model is tested on the held-out portion. The details can differ. In some cases there is a three way split, and in other a two way split. Some use k-fold cross validation while others use repeated random cross validation.

The main point of doing this type of validation is to help find the model that best balances the bias-variance tradeoff to give the most accurate predictions possible in future applications of the model. This bias-variance trade-off is critical in predictive modeling. Prediction error in future applications of the model is comprised of three components – bias, variance, and irreducible randomness. Too simple of model will have high bias but low variance in future applications. In the extreme, consider an intercept-only linear regression model. It will have zero variance in future predictions but high bias. At the other extreme, an over-fit version of highly complex model such as a principal pursuit regression or neural network may have very low bias but very high variance, also leading to poor prediction accuracy in the future. The “goldilocks” point is where the sum of bias and variance is minimized for future applications of the model. However, we do not know what model corresponds to this point. Holdout testing is the main approach we have for trying to estimate where this point is.

Standard holdout validation approaches work well when the full data set is sufficiently large and diverse that the cases in the hold out samples represent future conditions. However, if the data set is not representative of future hazards, the results of hold-out testing can be very misleading. Consider our example of predicting water system damage due to earthquakes. Hold-out testing may suggest that the out of sample prediction errors are small enough to be acceptable. However, this is really a *conditional* statement. The error estimates are conditioned on the data available for training and testing. If a future event is substantially different from what is in the data used to train and test the AI model(s), the errors may be much higher, making the hold out test results misleading about actual future model performance. This is particularly critical if the training set does not contain high-intensity hazard events. In this case, strong holdout results can be very misleading if the model is later applied to a stronger event. *Great care must be taken in interpreting the results of holdout testing as they can be misleading if the training and testing data is not representative of future conditions*.

A particular variation of the holdout testing challenge that requires particular attention is the case of zero-inflated data. Zero-inflation occurs in a data set when there are substantially more zeros, no damage events in a natural hazards setting, than what standard distributions and models can account for. If standard error metrics such as mean absolute error or mean squared error are used with highly zero-inflated data, the "best" model (by those metrics) is usually one that is heavily skewed towards predicting zero in nearly all settings. Indeed, the "best" model by these metrics in many zero-inflated cases is the model that always estimates zero impacts. This is useless in practice. Error metrics, model training, and holdout testing approaches must be adapted to handle zero-inflation. Error metrics that weight errors in the non-zero class higher (e.g., Shashaani et al. 2018) can help. Recent research has also found that multi-stage predictive models explicitly modeling the zero-class can offer improved predictive accuracy (e.g., Kabir et al. 2018). Traditional approaches for addressing unbalanced data focus on rebalancing approaches in which the training data is modified to be more balanced through undersampling the non-zero records, replicating some subset of the zero records, generating new zero-value records, or some combination of these. These approaches can be useful (e.g., Kabir et al. 2019), but the challenges with traditional error metrics driving model selection towards those biased towards predicting zero remain.

A third key challenge comes when the model results are conveyed to decision makers. It is difficult to convey model accuracy and the uncertainty that is inherent in any AI model output to decision makers in a way that they (1) will understand and (2) can use to improve decision making. Point estimates, while commonly provided, are misleading because they do not convey the often considerable degree of uncertainty in the predictions. At the same time, stating that a model is X% accurate (e.g., 95% accurate) is a meaningless statement unless the error metric on which the statement is based is clearly conveyed and agreed upon by the model users.

From a technical point of view, moving to probabilistic models (e.g., Kabir et al. 2019) is a significant improvement. Probabilistic models capture at least some of the uncertainty inherent in predictions, and this should improve the information basis for risk management decision making. However, we have found that it is problematic in practice to communicate the results of probabilistic models to decision makers in ways that they will actually understand. Consider, for example, a model that estimates the cumulative density function (CDF) of the number of person-hours needed for power restoration from a storm. We have found that providing the full CDF to decision makers often leads to confusion and that using summary measures such as confidence intervals leads to mis-interpretation and the loss of much of the probabilistic information. With some decision makers we have had success by communicating the probability of exceeding specified thresholds (e.g., P(required hours > 1000)). However, significant research gaps remain in the understanding of how practical decision makers interpret and understand different approaches for communicating uncertainty in predictions.

Overall, great care must be taken in communicating the results of an AI model so that those using the model understand what the results really mean and not be misled into a false sense of precision and certainty in the predictions.

6. Moving Forward with AI for Natural Hazard Risk Analysis

AI methods are of critical importance for natural hazards risk analysis, and the field will continue develop and use these methods. However, we must take care that we avoid the pitfalls associated with the use of AI methods for natural hazards risk analysis. These problems, discussed above, can lead to false confidence in model output and mislead decision makers. However, not using AI methods deprives decision makers of an important set of tools for better leveraging large, complex data sets to gain enhanced insights and improve their decision making.

While AI methods are already important for natural hazards risk analysis in both practice and research, more research is needed to improve AI methods and how they are used in risk analysis. Better methods for validating models and stress-testing models in a way that better aligns with the requirements of natural hazards risk analysis are needed. In particular, better methods are needed for situation in which the training data does not contain events as strong as those for which the model will be used in the future. We also need to continue to develop improved methods for modeling based on zero-inflated data. Due to the nature of the events that we model as risk analysts, zero-inflation is common in natural hazards risk analysis. We also need to develop better ways of communicating the results of models, including their limitations and uncertainties, so that decision makers better understand and appreciate both the predictions and the limitations of these predictions. And we need to ensure that predictive models used to support natural hazard risk analysis are subjected to rigorous, independent peer review to maintain high standards of technical rigor.

AI methods have an important role to play in natural hazards risk analysis, and risk analysis researchers must help develop these methods and their use in a way that best meets the challenges of our field.

Acknowledgements

This work was partially funded by a grant from the U.S. National Science Foundation (grant number 1638197). The views in this paper are those of the author and do not necessarily represent those of the National Science Foundation.

References

Aven, T. (2013). A conceptual framework for linking risk and the elements of the data–information–knowledge–wisdom (DIKW) hierarchy. *Reliability Engineering & System Safety*, *111*, 30-36.

Aven, T. (2017). Improving risk characterisations in practical situations by highlighting knowledge aspects, with applications to risk matrices. *Reliability Engineering & System Safety*, *167*, 42-48.

Badr, H. S., Zaitchik, B. F., & Guikema, S. D. (2014). Application of statistical models to the prediction of seasonal rainfall anomalies over the Sahel. *Journal of Applied meteorology and climatology*, *53*(3), 614-636.

Baroud, H., & Barker, K. (2018). A Bayesian kernel approach to modeling resilience-based network component importance. *Reliability Engineering & System Safety*, *170*, 10-19.

Burton, S., Gauerhof, L., & Heinzemann, C. (2017, September). Making the case for safety of machine learning in highly automated driving. In *International Conference on Computer Safety, Reliability, and Security* (pp. 5-16). Springer, Cham.

Chan, P. K., & Stolfo, S. J. (1998, August). Toward Scalable Learning with Non-Uniform Class and Cost Distributions: A Case Study in Credit Card Fraud Detection. In *KDD* (Vol. 1998, pp. 164-168).

Dueñas-Osorio, L., & Vemuru, S. M. (2009). Cascading failures in complex infrastructure systems. *Structural safety*, *31*(2), 157-167.

Francis, R. A., Guikema, S. D., & Henneman, L. (2014). Bayesian belief networks for predicting drinking water distribution system pipe breaks. *Reliability Engineering & System Safety*, *130*, 1-11.

Guikema, S. D. (2009). Natural disaster risk analysis for critical infrastructure systems: An approach based on statistical learning theory. *Reliability Engineering & System Safety*, *94*(4), 855-860.

Guikema, S. D., Nateghi, R., Quiring, S. M., Staid, A., Reilly, A. C., & Gao, M. (2014). Predicting hurricane power outages to support storm response planning. *IEEE Access*, *2*, 1364-1373.

Han, S. R., Guikema, S. D., Quiring, S. M., Lee, K. H., Rosowsky, D., & Davidson, R. A. (2009). Estimating the spatial distribution of power outages during hurricanes in the Gulf coast region. *Reliability Engineering & System Safety*, *94*(2), 199-210.

Huang, Z., Yu, Y., Gu, J., & Liu, H. (2016). An efficient method for traffic sign recognition based on extreme learning machine. *IEEE transactions on cybernetics*, *47*(4), 920-933.

Kabir, E., Guikema, S. D., & Quiring, S. M. (2019). Predicting Thunderstorm-Induced Power Outages to Support Utility Restoration. *IEEE Transactions on Power Systems*, *34*(6), 4370-4381.

Kaplan, S., & Garrick, B. J. (1981). On the quantitative definition of risk. *Risk analysis*, *1*(1), 11-27.

Kircher, C. A., Whitman, R. V., & Holmes, W. T. (2006). HAZUS earthquake loss estimation methods. *Natural Hazards Review*, *7*(2), 45-59.

LaRocca, S., Johansson, J., Hassel, H., & Guikema, S. (2015). Topological performance measures as surrogates for physical flow models for risk and vulnerability analysis for electric power systems. *Risk Analysis*, *35*(4), 608-623.

Liong, S. Y., & Sivapragasam, C. (2002). Flood stage forecasting with support vector machines 1. *JAWRA Journal of the American Water Resources Association*, *38*(1), 173-186.

Maes, S., Tuyls, K., Vanschoenwinkel, B., & Manderick, B. (2002, January). Credit card fraud detection using Bayesian and neural networks. In *Proceedings of the 1st international naiso congress on neuro fuzzy technologies* (pp. 261-270).

Mosavi, A., Ozturk, P., & Chau, K. W. (2018). Flood prediction using machine learning models: Literature review. *Water*, *10*(11), 1536.

Pomerleau, D., & Jochem, T. (1996). Rapidly adapting machine vision for automated vehicle steering. *IEEE expert*, *11*(2), 19-27.

Quiring, S. M., Schumacher, A. B., & Guikema, S. D. (2014). Incorporating hurricane forecast uncertainty into a decision-support application for power outage modeling. *Bulletin of the American Meteorological Society*, *95*(1), 47-58.

Rivero-Calle, S., Gnanadesikan, A., Del Castillo, C. E., Balch, W. M., & Guikema, S. D. (2015). Multidecadal increase in North Atlantic coccolithophores and the potential role of rising CO2. *Science*, *350*(6267), 1533-1537.

Schneider, P. J., & Schauer, B. A. (2006). HAZUS—its development and its future. *Natural Hazards Review*, *7*(2), 40-44.

Shashaani, S., Guikema, S. D., Zhai, C., Pino, J. V., & Quiring, S. M. (2018). Multi-Stage Prediction for Zero-Inflated Hurricane Induced Power Outages. *IEEE Access*, *6*, 62432-62449.

SRA 2018. *Society for Risk Analysis Glossary*, updated August 2018. Available at https://sra.org/sites/default/files/pdf/SRA%20Glossary%20-%20FINAL.pdf. Last accessed August 23, 2019.

Suryanita, R., & Adnan, A. (2012). Intelligent Monitoring System on Prediction of Building Damage Index using Neural-Network. *Telkomnika*, *10*(1).

Tripathi, S., Srinivas, V. V., & Nanjundiah, R. S. (2006). Downscaling of precipitation for climate change scenarios: a support vector machine approach. *Journal of hydrology*, *330*(3-4), 621-640.

Wang, S., Hong, L., Ouyang, M., Zhang, J., & Chen, X. (2013). Vulnerability analysis of interdependent infrastructure systems under edge attack strategies. *Safety science*, *51*(1), 328-337.

Winkler, J., Duenas-Osorio, L., Stein, R., & Subramanian, D. (2010). Performance assessment of topologically diverse power systems subjected to hurricane events. *Reliability Engineering & System Safety*, *95*(4), 323-336.