Machine learning
the role of machines for resilient communities

Kammouh, Omar; Cimellaro, Gian Paolo

DOI
10.1061/9780784415894.ch5

Publication date
2022

Document Version
Accepted author manuscript

Published in
Objective Resilience

Machine Learning: The Role Of Machines for Resilient Communities

Omar Kammouh[a] and Gian Paolo Cimellaro[b]

[a] Department of Materials, Mechanics, Management & Design (3Md)

Delft University of Technology

Stevinweg 1, 2628 CN, Delft, Netherlands,

o.kammouh@tudelft.nl

[b] Department of Structural, Geotechnical and Building Engineering (DISEG)

Politecnico di Torino

Corso Duca degli Abruzzi 24, 10129, Torino, Italy

gianpaolo.cimellaro@polito.it

Abstract

Smartphones, wearables, and internet-connected devices are constantly with us. They generate an enormous amount of information with different formats and unrelated meanings. A human would be able to find links between the information and conclude if an emergency is going on. Nevertheless, the human capacity to process information is limited to small sets of data. This chapter introduces the role of Machine Learning (ML) in Resilience Engineering and discusses actual cases of emergencies where ML contributed positively. Different examples of emergency scenarios, from natural to manmade hazards, are presented and the contribution of ML is highlighted. The limitation of ML due to data scarcity is equally important and also discussed.

This chapter encourages practitioners to integrate ML techniques and Artificial Intelligence (AI) into their emergency plans. It invites them to organize programs that aim at training their researchers and employees to use AI to deliver their job. One key aspect that the chapter stresses is that machines can never act on their own without a modicum of human involvement.

1. Introduction

Of all the dreams of humankind, the most popular one is certainly the ability to predict the future. By staring at a crystal ball or the stars, different people in the past have developed different techniques to fight the scariest of all potential demons - uncertainty. They may have done this for one simple reason which is knowing in advance what is going to happen. Unfortunately, that is not always the case in practice. Take an example of the slopes of Vesuvius which currently host the homes of 3 million inhabitants. Even though science has been very clear that a new explosive eruption will occur sooner or later (Barnes 2011), people still live there. A similar situation exists at the Campi Flegrei caldera. (Kilburn et al. 2017).

1.1 Resilience Definition

In the context of this chapter, resilience is the ability to withstand stresses caused by external events and recover quickly to the functional state (Kammouh et al. 2018). Resilience ensures a reliable and affordable continuity of the service supply in normal operation as well as during (and after) disaster events. Several methods to quantify the resilience of communities exist in the literature (Cimellaro et al. 2016; Kammouh et al. 2017; Kammouh et al. 2018; Kammouh et al. 2019). However, none has considered the role of Machine Learning (ML) in their respective assessments of resilience.

According to Bruneau et al. (2003), the resilience of a system depends on its functionality performance. The functionality of a system is the ability to use it at an impaired level. The conceptual approach of resilience described in (Bruneau et al. 2003) is illustrated in Figure 1. The functionality performance (Q) ranges from 0 % to 100 %, where 100% and 0% imply full availability and non-availability of services, respectively. The occurrence of a disaster at time t0 causes damage to the system and this produces an instant drop in the system's functionality (ΔQ). Afterward, the system is restored to its initial state over the recovery period (t_1-t_0). The loss in resilience is considered equivalent to the quality degradation of the system over the recovery period. Mathematically, it is defined by Eq. (1):

$$LOR = \int_{t_0}^{t_1} \left[100 - Q(t)\right] dt \tag{1}$$

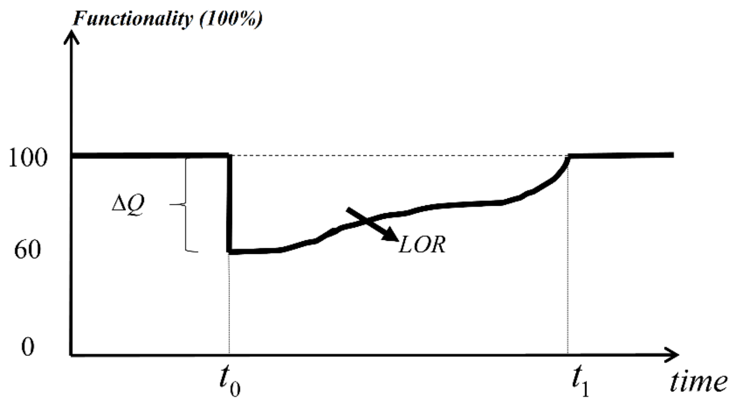

Figure 1: The concept of Disaster Resilience

where LOR is the loss-in-resilience measure, t_0 is the time at which a disastrous event occurs, t_1 is the time at which the system recovers to 100% of its initial functionality, Q(t) is the functionality of the system at a given time t.

1.2 Machine Learning and Artificial Intelligence

Artificial Intelligence (AI), and its subset Machine Learning, have the potential to offer valuable solutions to achieve resilient communities. ML is employed in a range of computing tasks where designing and programming explicit algorithms with good performance are difficult or infeasible. To understand its benefits within the resilience-relevant aspects (social, economic, infrastructural, institutional, environmental, and community-wise), the role of ML in the different disaster management applications is discussed:

1. Model identification: ML can learn patterns and provide indicators for future predictions. This is what researchers are constantly trying to do with natural and human disasters. In fact, ML is much better than humans at learning from mistakes, literally.
2. Emergency detection: in emergencies, choosing one alternative over another can cost lives and money. Questions like "Which building needs to be addressed first?" or "Is it safe to send the civil defense in this area?" need precise and quick answers. ML can detect if something unusual is happening, trigger intelligent alerts, and suggest the optimal ways to deal with the emergency.
3. Solution generation: expecting constraints and requirements as input, AI techniques explore the entire solution space, then investigate every solution that may solve the problem.

1.3 Semantic Representation of Emergency

Thanks to the internet, we are all connected. We are given an easy way to share multimedia content in real-time, making it available not just to our chosen emergency contact but to a whole audience. Smartphones, wearables, and the Internet of Things (IoT) devices are constantly with us: they save our location, our pictures, our voices. All this generates an enormous amount of information, in very different formats, with very different and unrelated meanings. While humans are capable of understanding and using this information to figure out if there is an emergency going on, machines are much more efficient in performing such a task considering that several emergencies are taking place at the same time. Nonetheless, a machine would struggle more to find meaning in the data.

Hence, at the heart of any ML approach to emergencies is the representation of the real-world data, in a language that is comprehensible to the machines. The Semantic Web (SW) is a set of technologies that provide standardized formats for the representation of both data and ontological background knowledge (Tresp et al. 2006). Here, by ontology, we mean the domain-specific background information organized in logical statements. An ontology describes object classes, predicate classes, and their interdependencies. Using this common vocabulary, machines communicate and understand. This is exactly what is happening in the background when we type on Google "Brad Pitt’s mother". First of all, it understands our question. Then, it starts exploring the Google Knowledge Graph, a graph where every edge is a relationship between two entities (in this case, Brad Pitt, and his mother), to extract the answer to our question. Google is not just listing top articles containing the same words we have inserted in our query: it is instead producing an intelligent answer because it has really understood our question.

Ontologies are built on top of two standards: RDF and RDFS. RDF is a resource description framework which represents information about resources using basic triplets: subject, predicate, object. Each resource is associated with one or several concepts (i.e., classes) via the type-property. Concepts are defined in the RDF Vocabulary Description Language, also called RDF-Schema or RDFS. The web ontology language is OWL, which allows stating that classes are equivalent or disjointed and that properties and instances are identical or different. Properties can be symmetric, transitive, functional, or inverse functional. In RDFS concepts are simply named, while OWL allows the user to construct classes by enumerating their content (explicitly stating its members) or by forming intersections, unions, and complements of other classes. An ontology formulates logical statements, which can be used for analyzing data consistency and for deriving new implicit statements concerning instances and concepts.

So, what does ML have to do with all this? ML comes into play with ontology evaluation, refinement, evolution, as well as the merging and alignment of ontologies (Tresp et al. 2006). One possible scenario is the following: we build an ontology, a representation of the world which becomes our baseline. Using ML, we can apply learning algorithms to our axioms and instances, which in turn allows us to understand more about our world.

We can extract new subject-predicate-object triplets, that will then be added to our ontology to generate more knowledge. ML would then need to create samples of the population existing in the ontology and extract the latent features (the fundamental characteristics) introduced in a cluster analysis or a principal component analysis (PCA), with the support of SQL (declarative querying language) or SPARK (big data framework). Finally, ML would generate new statements which would be weighted depending on their likelihood: after all, ML still lives in the dimension of the uncertain. This likelihood can be established by ensemble methods: different algorithms with different characteristics and different results that are merged to form a more likely result.

ML can also be employed in ontology learning. This includes the identification of concepts, concept hierarchies, properties, property hierarchies, domain, and class definitions. One way to do this is by applying hierarchical clustering techniques like single-link, complete-link or average-link clustering to leverage the semantic and syntactic context of words to understand new concepts previously absent from the ontology (Tresp et al. 2008). This idea has been applied to build a crowdsourcing-based knowledge base, that is extracted from social media keywords and patterns (Xu et al. 2016).

To sum up, ML is fundamental in times of crisis and emergency management because it provides an underlying dictionary that allows us to understand what is happening, how to react, how to communicate with different systems to dispatch alerts. It is also a way to incorporate the new knowledge from the data and represent it in a formal way that makes it available not just to a single script but to entire systems. Starting from a baseline that comes from theory (a theoretical, physical model created by earthquake experts), it can then add more knowledge extracted from the data.

2. Model Identification

"*All models are wrong, some are useful*" - G. Box

In 2014, it was estimated that natural and man-made catastrophes took 7700 lives and caused approximately $ 110 (US) billion in damages. The need (and the market potential) for predictive tools is extremely clear.

2.1 The Problem of Data Integration

In data science, a simple algorithm with a lot of data is considered to be better than a complex algorithm with far less data. Very often in ML and data science, the fundamental problem is the lack of data. And by data, we don't mean just any kind of data, but rather meaningful, labeled, organized data that can be used consistently by any algorithm. As mentioned earlier, our world is becoming more and more connected. The internet of things is the term used in the tech community to describe the existence and communication of different sensors and devices through the internet. One example that leverages this task force of measuring sensors is the Quake-Catcher Network (QCN), a joint seismic initiative that has provided traditional seismic stations with innovative data sources, bringing together information from the accelerometers in mobile phones and cloud computing and guaranteeing faster detection of earthquakes. This stems from a very democratic, crowd-sourcing idea: everybody can contribute to providing better-performing emergency response systems at a low cost (Cochran et al. 2009).

A key to more data and more accurate results is often the integration of multiple sources. One model was able to detect landslides using a Bayesian approach using social and physical sensors, such as USGS seismometers and TRMM satellites (Musaev et al. 2014). The system periodically downloads data from multiple social and physical sensors, extracts information from social sensors like Twitter, YouTube, and Instagram, then performs multiple filtering steps, of exclusive or inclusive type. These filters were related to the specific type of the emergency: based on sentinel words or phrases, geo-tags, an ML classification component, a blacklist of URLs. The result of this filtering is merged with the one coming from physical sensors, such as seismic activity or rainfall levels measurements. These steps are included in Figure 2 for further clarity.

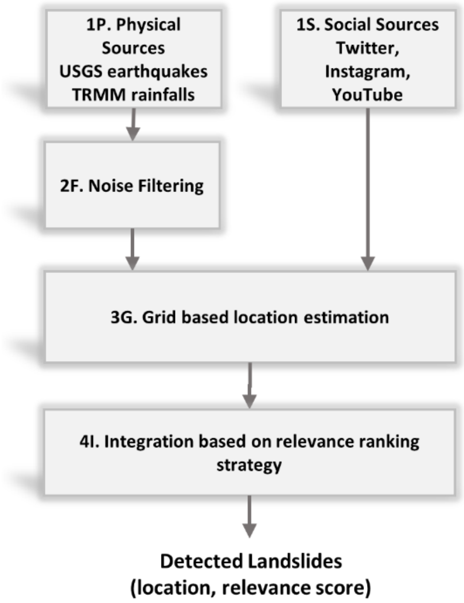

Figure 2: Overview of the data flow for a landslide's detection application

Big Data-enabled integration was also the fosterer of a flood-detection system. Researchers combined information from Twitter and from Satellite observations, to build a learning and real-time map of floods. The problem of integration is also behind Digital Delta, a research program involving IBM, the Rijkswaterstaat, the University of Delft, and the Deltares Water Institute (Byrne). It has proven that by listening to what the data have to say, it is possible to build better infrastructure, understand the weakest points of the current infrastructure, and achieve better target maintenance and investments. However, this is not just a matter of data integration, it is also a matter of response integration amongst the many districts and communities.

2.2 Predicting Natural and Man-made Hazards

We have been supported by AI in various fields. Now, researchers have found that AI can be used for natural disaster prediction. AI can forecast the occurrence of multiple natural disasters given large good-quality datasets. Examples of the natural hazards predictable for AI are earthquakes, volcanic eruptions, and hurricanes.

Many seismic scholars and scientists believe that predicting earthquakes is nearly impossible. But thanks to new model identification and ML techniques, a lot of interesting insights are being extracted from seismic data. Researchers are using deep learning systems to gather large quantities of seismic data for analysis (Zhang et al. 2018). AI may use seismic data to evaluate earthquake magnitude and frequency. These data can be useful in forecasting the occurrence of earthquakes. Some attempts have shown that AI-based algorithms can predict aftershock positions more precisely than other approaches.

Volcanos eruptions prediction has always been a challenge. Recent attempts could find ways of accurately forecasting volcanic eruptions by training an AI system to recognize tiny volcanic ash particles. The ash particle shape can be used to classify the volcano's type. These advances can help to predict eruptions and to establish strategies for minimizing volcanic hazards.

Hurricanes are one of the most damaging natural hazards. NASA recently employed a system that combined satellite images and ML to monitor Hurricane Harvey. The system proved to be six times better than the conventional monitoring systems: the hurricane can be monitored every hour instead of every six hours as in the case of traditional systems. Therefore, technical advances are helping to track hurricanes and forecast the course of hurricanes which can aid in mitigation efforts.

For man-made hazards such as terrorism, it is reasonable to express doubt with a question such as: is there really nothing we can do to prevent, if not predict, terrorism? In the aftermath of a terror attack, much controversy is sparked when it turns out that the terrorist organizations were very well "known" to authorities. But what seems to be the key issue is that it is extremely difficult for governing entities to track every single individual who has demonstrated a weird or dangerous behavior that would lead to terroist-like behavior. This is where ML could be of use: it is not only a matter of automating and repeating a task (that of monitoring an individual), which is something machines can do very well. What is needed is continuous monitoring of a number of different sources and combining them into one, meaningful output. Again, as mentioned earlier in this chapter, it is always a matter of humans and machines: human intervention will always be required in the end to extract a decision from all this information. But this will be an informed decision, an educated and science-backed guess.

Some researchers have already tried to experiment with the potential that lies in the application of ML techniques for emergency detection (Tutun et al. 2017). The researchers attempted to identify patterns in suicide attacks using ESALLOR, a new Evolution Stimulating Annealing Lasso Logistic Regression. The system identified the most important features of terror attacks, while also proposing a new similarity function to estimate the relationship among similar events.

Machine-learning classifiers are in general very good at discovering trends, clusters, and stereotypes. They are statistical approaches, not individualistic. While it is ok for a recommender system like Spotify to suggest a song you don't really like just because other users, that have proven in general to have a musical taste similar to yours, enjoyed them before, it is less okay for the government to increase surveillance on you, intercept your communications, monitor what you do (a violation of constitutional rights and a waste of law enforcement resources).

ML could also recognize faces via ordinary monitoring systems (CCTV). The FBI, for instance, has access to nearly 412 million photos in its facial recognition system (Orcutt 2016), which constitutes a great training set for learning algorithms. State-of-the-art face matching systems can be nearly 95 percent accurate on mugshot databases which sounds extremely promising, but these pictures are very clear and taken in controlled environmental conditions and of cooperative subjects. Adding blurred, dark pictures may be characterized by unusual facial expressions or poses, which would worsen the accuracy. Moreover, any gender, age group, or race that is under-represented in the training data will be reflected in the algorithm performance. This is probably the reason why some organizations that are using MorphoTrus's facial and iris recognition are still uncertain the accuracy of the system.

In the absence of faces, ML could also identify terrorists from their victory sign, using hand shape biometrics (hand silhouette, finger widths, lengths, angles, etc.). Image segmentation is an important processing step in many images, video, and computer vision applications, and it was the key to the victory sign analysis. In this chapter, we mention four approaches to segment the hand: Otsu's method of histogram shape-based image thresholding (Xu et al. 2011); K Nearest Neighbors classifiers that distinguish between "hand" and "not hand" using Euclidean (Laaksonen and Oja 1996); Manhattan and Hassanat distance (Alkasassbeh et al. 2015), and

Artificial Neural Network (ANN) based on RGB information (Ramil et al. 2018). The training architecture is shown in Figure 3.

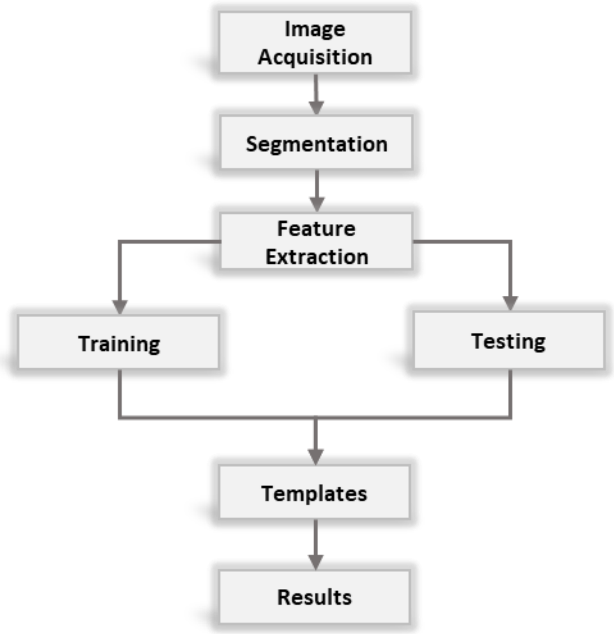

Figure 3: A typical hand shape biometric system

Given the above, ML has clear advantages: it easily identifies trends and patterns, no human intervention is needed, etc. However, it also has disadvantages due to data acquisition issues, time and resource requirements , data interpretation difficulties, etc.

3. Emergency Detection

3.1 Detecting and Managing the Emergency

During emergencies, it is of utmost importance to be able to understand where the emergency is and what has been damaged the most. In the case of particularly big emergencies, it is even harder to be able to organize the available human and financial resources. Many advancements in recent technologies have been useful to partially tackle this problem.

One of the first studies on this topic developed an ML tool predicting the damage expected on a network based on the weather forecast (Angalakudati et al. 2014). In particular, it had in mind what today we would call "Industry 4.0", where many sensors work together creating a robust monitoring system that helps prevent a million-dollars-system's failure. If we think of an electrical network, weather-related damage might result in a huge economic loss where several days are needed to restore the situation back to normal conditions.

It is reasonable to imagine a feature where drones fly above a critical location in real-time, or where heat-detecting robots are able to locate survivors and perform rescue operations more quickly and efficiently than a team of humans are capable of doing. Embedded systems and IoT applications are going to be our eyes and ears across the world, providing more and more accurate information concerning people and buildings. This allows better planning from the rescuers part, which can have a clear idea about the topography of the landscape and the extent of damage to a building.

When an emergency occurs, two approaches can be utilized to gain further information. First, it can be detected from the real world itself, thanks to the ubiquitous presence of sensors throughout the world. Second, we can rely on the immediacy of social networks and news agency reporting. Both are these approaches are discussed below.

3.2 Emergency Detection: Real World

Traditional warning systems operate in a broadcast fashion (Cipolla et al. 2016). Sirens, text messages, or emails are meant to alert almost everyone, in every place, and every situation. Cellular phone or radio broadcast networks make it hard for these systems to reach individuals who are located inside buildings. Moreover, networks such as Ethernet and WiFi tend to fail in times of extremely high demand (like emergencies). In these situations, deep learning can be used to trigger emergency warning systems via existing infrastructure such as closed-circuit television (Kang and Choo 2016). This approach is to start from a real-time video analysis: CCTV modules store the captured video data locally and periodically monitor the footage received performing object detection and image classification. When an emergency is detected, an alert is forwarded directly to the police station. This way, emergency detection is autonomous, and civil protection receives more and more accurate information about the emergency (e.g., type, location, time, images, etc.) The two types of emergencies aforementioned are generated via a Poisson process, progressively increasing the level of strength (weak, normal, strong) and the lambda value. This deep learning approach makes the overall system more scalable and faster, as it can be directly deployed in embedded devices (such as CCTV) and respond extremely quickly (in milliseconds). Deep learning also guarantees that no features need to be hardcoded by experts as they will be learned by the network.

There have been several attempts to use ML is early warning systems to predict natural disasters and processes. For instance, Asnaning and Putra (2018) introduced the automatic water level recorder (AWLR) in conducting water level monitoring at the water-gate dam. The function of AWLR sensor is for monitoring and recording in a database with real-time sensing. The results show that the low-cost AWLR sensor has reduced processing time by 92.7% compared to conventional data processing. Another is applying ML to an early warning system for very short-term heavy rainfall (Moon et al. 2019). The authors introduced a method for an effective early warning system for very short-term heavy rainfall with ML techniques. Results showed a better predicting pattern than other methods (Moon et al. 2019).

3.3 Emergency Detection: Virtual World

Social networks and internet platforms, in general, have been hosting people's messages and thoughts for quite some time now. Often, these messages have been frequently analyzed using simple techniques, such as measuring the frequency of emergency related words as the emergency is approaching. These messages are real-time, can be location-based, and ultimately provide precious information about disease outbreaks (Brownstein et al. 2007), conflicts and terror-related situations, and natural catastrophes. We can see this very clearly from the Boston Marathon terrorist incident (Cassa et al. 2013).

Twitter, among others, is a very valuable source of information. From one side, it carries precious and real-time insight into events as they evolve. On the other side, care must be taken to avoid false-positive reports with negative effects. For this reason, it is necessary to compare the cost of unnecessary investigation and the opportunity cost of not reacting early enough (Corvey et al. 2010).

A traditional approach in natural language processing is the Bag Of Words model (Araque et al. 2017), where a document is mapped to a feature vector, and then classified by ML techniques. This is a very simple approach, and it destroys information like word order and syntactic structures. Another kind of feature that can be used is Part Of Speech (POS) tagging, which is commonly used during a syntactic analysis process. Some authors refer to this kind of feature as surface forms, as they consist of lexical and syntactical information that relies on the pattern of the text, rather than on its semantic aspect. These low-level classifiers can be used in rule-based approaches, meaning that the low-level predictions are treated by rules such as majority voting, or in meta-learning, where they constitute features (parameters) for higher-level models.

Combining classifiers usually achieves greater accuracy and single classifiers alone. This integration can happen concurrently (divide the original dataset into several subsets from which multiple classifiers learn in a parallel fashion) as it happens in bagging, or sequentially, such as boosting. In Natural Language Processing, deep learning has been used to learn word vector representations using neural language models such as word2vec (Collobert et al. 2011). This approach models words as vectors, allowing them to retain a huge amount of syntactic and semantic regularities.

Unsupervised learning has also been employed, for example via autoencoder, which allows extracting a new, more concise (or de-noised) representation of the input. In general, there is a growing tendency that tries to incorporate additional information to the word embeddings created by deep learning networks. Augmenting knowledge in the embedding vectors with other sources of information can also be useful, for example using a previous related topic or sentiment related information.

A very recent work proposes the Recursive Neural Tensor Network (RNTN) model (Araque et al. 2017), which represents a phrase using word vectors obtained in an unsupervised manner and a parse tree, computing vectors for higher nodes in the tree using a tensor-based composition function. On top of this, there is the ensemble model which combines classifiers trained with deep and surface features. This model combines several base classifiers into one ensemble that makes predictions from the same input data. This model is proposed to combine several types of features into a unified feature set and, consequently, combine the information these features give. In this way, a learning model that learns from this unified set could achieve better performance scores that one that learns from a feature subset.

The Qatar Computing Research Institute (QCRI) has developed a free, open-source, ML-based framework to improve efficiency and management in the aftermath of crises: AI for Disaster Response (AIDR) (Imran et al. 2014). Its objective is to help create a comprehensive picture of an emergency, helping the organization of the emergency operation centers. According to tweet analysis, the system can identify and categorize needs based on urgency, infrastructure damage, and resource deployment. The rescuers can reduce the time spent on planning and organization and can focus instead on helping those who need help. Organized reaction and targeted alerting (contacting the people in the identified places) can help evacuate people quickly from the identified danger zones.

3.4 Managing the Emergency

Once the emergency is detected, then a planned intervention is to be deployed. Several companies are already involved in this field, experimenting with several learning-based solutions. One example is IBM, which has developed a predictive tool, the "Intelligent Operations Center for Emergency Management", in partnership with the Weather Channel. The system integrates multiple data sources in real-time to create "multifaceted situational awareness of city resources & events and create a collaborative environment for planning, monitoring & sharing information”.

The information retrieved from this kind of analysis can be very useful in the planning of the evacuation or rescue activities after an emergency or crisis. It could be for example included in models such as a dynamic Bayesian network (DBN) (Radianti et al. 2015), supporting distinct kinds of crowd evacuation behavior, both descriptive and normative (optimal). Descriptive modeling is based on studies of physical fire models, crowd psychology models, and corresponding flow models, while we identify optimal behavior using Ant-Based Colony Optimization (ACO). Simulation results demonstrate that the DNB model allows us to track and forecast the movement of people until they escape, as the hazard develops from time step to time step. Furthermore, the ACO provides safe paths, dynamically responding to current threats, such as cyber threats (Kammouh and Cimellaro 2018). This kind of model integrates concepts from graph theory and probability theory, capturing conditional independencies between a set of random variables by means of a directed acyclic graph (DAG), each edge of which typically represents a cause-effect relationship.

A similar path is being followed by One Concern, a machine-learning-based startup that provides emergency operations centers (EOCs) with critical situation awareness; for instance, instant information on response priorities and other insights to allocate all the limited resources effectively. The platform sends automatic alerts when an earthquake seems to have affected a certain county, including key information like "the elderly population in a particular block that is badly damaged, or the number of kids in a school which could be hit" (Shueh 2016). The system can also ease the creation of the Initial Damage Estimate (IDE), being able to identify and quantify the extent of damage to his jurisdiction with a significant amount of accuracy in minutes, thus saving a lot of time, and promising high precision. The system puts special care in redundancy and distributed servers, allowing the platform to be up and running even when phone networks are usually down (indeed, during crises).

Concerning the technology used, very little is known because it is proprietary. What is known is that the same technology used for real-time estimation is also included in an AI-based training module that will allow emergency operations centers to train on scenarios based on actual simulations to get a real sense of the situation, helping personnel readiness and plan development, thereby making a community more resilient.

4. Solution Generation and Decision Making

4.1 An Excursus on AI

The key aspect of every disaster management situation is what happens after the moment of solution generation. That is the moment when the emergency has gone, we have counted the injuries and the victims, we have calculated losses and damages and it is now time to build again (De Iuliis et al. 2019; Kammouh et al. 2018). History has shown that sometimes this second chance is not well-used. This field has great potential for AI and ML applications. The history of computer science leads us to imagine enormous supercomputers producing the result of very complex, yet mechanical calculations. Amongst all its qualities, we would certainly not define a machine innovative (Perez. 2016). Surprisingly enough, a new branch of AI research is producing generative design tools, algorithms that ask for four ingredients: goals, constraints, computing power, and time. In return, they produce solutions that humans could have never come up with. How? They simply start from scratch and then, very methodically, search the entire solution space and explore every single possibility that fulfills the initial requirements (Conti 2017). The three structural elements shown in Figure 4 are all designed to carry the same structural loads and forces. As we move from left to right, we shift from a traditional design to the most recent computer optimization.

Figure 4: Evolution of a structural element using ML and computer optimization. Modified from (Carlos 2016)

With respect to the traditional production methods, generative solutions offer a height reduction of 50 %, weight reduction per node is 75%, and an overall weight reduction (on a construction project) of more than 40% (Carlos 2016). In this case, the strength of the machine over the human is that it is not biased: when the search algorithm starts, it is still a kid. It has no ideas about what has been studied for centuries, what is already working well, what has already been tested useless. It analyses every single possibility, without prejudice.

This problem cannot be tackled by "ordinary" Machine Learning. As we have seen so far, ML is the art of extracting the most important features from the data since it was designed to operate on known objects, not to invent them. Independently from the specific algorithm, learning problems usually look for a function that is a good representation of the mapping between objects and their corresponding classes. Learning models are not designed to hypothesize about the creation of new objects, they simply assume that by applying a series of operations we can learn new knowledge from that world by generalizing upon existing objects or relationships. These algorithms thus neglect the fact that sometimes it is simply more important to decide what to look for then finding what is already there. By contrast to decision and learning paradigms, the design is the creation of new objects. Designers generate multiple novel object definitions that might be explored next. The true value of a designer lies in their judgment. It is not a matter of choosing the best among existing objects, but to explore among a set of novel definitions. This is a decision theory specific to design processes, that is yet to be formulated (Kazakci 2014).

4.2 Resilience and the Role of Machine Learning

The impact of ML on the aftermath of an emergency is extremely relevant also from another point of view: we can image a central ML engine that considers all the most relevant variables like weather/geologic conditions, human exploitation, civil use of the building, history (what emergencies happened there, what went wrong), and builds an eternal knowledge base out of them. It is not hard to envision how the PEOPLES framework (Cimellaro et al. 2016) could contribute to this, and take strong advance from such knowledge. This knowledge would not get lost with time, politics, or just a change in the team or the company that is in charge of the reconstruction. ML is a form of intelligence that continues to grow and becomes more accurate and comprehensive as time (and data available) accumulates. Once more, a semantic way of dealing with Big Data is fundamental.

Moving on to the act of reconstruction itself, an intelligent machine could coordinate the workers, incorporate vision and change the path and the project as it goes on and as new impediments arise, as new data becomes available. Machines would be thus greatly contributing to the resilience of our new cities and buildings, in their capability to "sustain a level of functionality or performance for a given building, bridge, lifeline networks, or community, over a period defined as the control time" (Cimellaro et al. 2010).

5. Discussion and Conclusions

This chapter introduced the role of Machine Learning (ML) in different applications and scenarios of Resilience Engineering, such as during natural and manmade disasters. Three main applications for ML in disaster management are discussed: Model-identification, Emergency detection, and Solution generation.

In the model identification, the problem of data scarcity is presented. Data needs to be complete before any meaningful results can be drawn. The solution to this is by improving the data type and increasing the data channels. In Emergency detection, the application of ML in different fields (e.g., physical, virtual) is highlighted. The role and objective of ML in every field can be very different. Finally, in the Solution generation section, the effectiveness of ML in supporting human with decision making is discussed. This was also supported by real examples where the machines could generate better solutions than the human.

5.1 On Human-Computer Interaction

There is one very famous scene in the movie, “I, Robot” 2004, one of the most famous modern movies about robots coming to life. That is when Will Smith finally unveils to the audience the origin of his long-living hatred towards machines. This dates back to his past when as a result of a severe car crash, two cars (including his) fell into a river. Together with the others, a little girl fell into the water with him. A robot came to rescue, but soon understood that a) he couldn’t save everyone and b) Will had a much higher chance at survival than the little girl. As a result, Will was saved, the child was not.

This brief but relevant scene leaves a lot of us wondering: is this the kind of world we are about the make come true? A world where the law of the jungle is going to prevail, and logic and formal rules are going to take the place of the emotions, comprehension, altruism? While it is very hard at this point to predict the course of research in AI, especially in emergency management, we would argue that for the time being, machines are given a goal to reach, they do not find their own. It is then a matter of the human beings behind them, the very ones that set the goals and the parameters to evaluate the success of an algorithm. Ultimately, it is a matter of those who write the basic rules the machines will have to respect.

Finally, if we think of an autonomous driving scenario, many people argue that they would rather be completely in charge of their vehicle. Think of an emergency: would you like to be in control of what’s going on, or would you trust an algorithm that somebody else has written? For us, humans, the most powerful beings on earth, it is hard to devolve authority to somebody else, giving up on our very own right to decide for ourselves. But if we think of it, for just one second, we will soon realize that we are not really in control of emergency conditions. Most likely, we act guided by fear, irrationalism, or anxiety. And we can make very, very stupid decisions. This is because at the very moment when we think it’s most important to be in control, we are not. Our decisions are the result of a random mixture of chance, the mood of the day, and past (biased) experience. Wouldn’t it be better if we could be guided instead by a machine that is not a victim of those evil antagonists but is instead able to remain vigilant in every situation and act for the best? Wouldn’t it be better if the world could come together and decide what are the rules the machines should obey and what are the success criteria every human should be satisfied with? It is of utmost importance to find an answer before we even forget we had a question.

Ultimately, it is a problem of understanding the deepest rules governing the human-computer interaction, which roles are going to become machine-based and which ones are going to be more and more human-based in the future. None of the approaches mentioned in this work could ever take place with only machines, nor only humans: all of them require the cooperation of the two parts, leveraging what each can do better. In emergencies, humans and machines have equally important roles.

5.2 Complex Decision Making Under Emergency Conditions

The key to better emergency management is better coordination between human and machine intelligence. ML can intervene and eventually free the human decision-maker from all the low-level analytical tasks and unleash their imagination and creativity to a level that machines themselves could never reach.

The power of ML lies in its ability to provide extremely valuable and meaningful information to the humans, and ultimately make a difference in the decision process. This information is extremely important, especially in emergency conditions, when life-or-death decisions are due in a matter of minutes. Provided that algorithms will continue to improve, and models will be more and more accurate, are humans ready to accept this power? Are they ready to include the results of ML into their decision-making processes, allowing them the same credibility they would allow to a trusted human advisor? Are humans ready to accept the inexorable, scientific results, and the huge transformations they would trigger on our society?

Thanks to our augmented capabilities, our world is going to change dramatically. We're going to have a world with more variety, more connectedness, more dynamism, more complexity, more adaptability, and, of course, more beauty. The shape of things to come will be unlike anything we've ever seen before. Why? Because what will be shaping those things is this new partnership between technology, nature, and humanity (Conti 2017).

6. Recommendations

- Good monitoring systems and meaningful data are the basis of effective machine learning systems. Thus, practitioners should first invest in building reliable monitoring systems.
- Training programs on Machine Learning should be arranged for the researchers in research institutes and IT employees in professional industries.
- Create programs that aim at coordinating human and machine intelligence for better results.

- Test the emergency system independently from Machine Learning to see the efficiency of employing machine intelligence.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Research Council under the Grant Agreement n° ERC_IDEAL RESCUE_637842 of the project IDEAL RESCUE—Integrated Design and Control of Sustainable Communities during Emergencies.

References

Alkasassbeh, M., Altarawneh, G. A., and Hassanat, A. (2015). "On enhancing the performance of nearest neighbour classifiers using hassanat distance metric." *arXiv preprint arXiv:1501.00687*.

Angalakudati, M., Calzada, J., Farias, V., Gonynor, J., Monsch, M., Papush, A., Perakis, G., Raad, N., Schein, J., and Warren, C. "Improving emergency storm planning using machine learning." *Proc., 2014 IEEE PES T&D Conference and Exposition*, IEEE, 1-6.

Araque, O., Corcuera-Platas, I., Sanchez-Rada, J. F., and Iglesias, C. A. (2017). "Enhancing deep learning sentiment analysis with ensemble techniques in social applications." *Expert Systems with Applications*, 77, 236-246.

Asnaning, A. R., and Putra, S. D. "Flood Early Warning System Using Cognitive Artificial Intelligence: The Design of AWLR Sensor." *Proc., 2018 International Conference on Information Technology Systems and Innovation (ICITSI)*, IEEE, 165-170.

Barnes, K. (2011). "Volcanology: Europe's ticking time bomb." *Nature News*, 473(7346), 140-141.

Brownstein, J. S., Freifeld, C., Reis, B., and Mandl, K. (2007). "HealthMap: Internet-based emerging infectious disease intelligence." *Infectious disease surveillance and detection: assessing the challenges—finding solutions. Washington: National Academy of Science*, 183-204.

Bruneau, M., Chang, S. E., Eguchi, R. T., Lee, G. C., O'Rourke, T. D., Reinhorn, A. M., Shinozuka, M., Tierney, K., Wallace, W. A., and Winterfeldt, D. v. (2003). "A Framework to Quantitatively Assess and Enhance the Seismic Resilience of Communities." *Earthquake Spectra*, 19(4), 733-752.

Carlos, P. (2016). "The Alien Style of Deep Learning Generative Design." <https://medium.com/intuitionmachine/the-alien-look-of-deep-learning-generative-design-5c5f871f7d10>.

Cassa, C. A., Chunara, R., Mandl, K., and Brownstein, J. S. (2013). "Twitter as a sentinel in emergency situations: lessons from the Boston marathon explosions." *PLoS currents*, 5.

Cimellaro, G. P., Reinhorn, A. M., and Bruneau, M. (2010). "Framework for analytical quantification of disaster resilience." *Engineering structures*, 32(11), 3639-3649.

Cimellaro, G. P., Renschler, C., Reinhorn, A. M., and Arendt, L. (2016). "PEOPLES: a framework for evaluating resilience." *Journal of Structural Engineering*, 142(10), 04016063.

Cimellaro, G. P., Zamani-Noori, A., Kammouh, O., Terzic, V., and Mahin, S. A. (2016). "Resilience of Critical Structures, Infrastructure, and Communities." Pacific Earthquake Engineering Research Center (PEER), Berkeley, California, pp. 318.

Cipolla, E., Rizzo, R., Stabile, D., and Vella, F. "A tool for Emergency Detection with Deep Learning Neural Networks." *Proc., KDWeb*.

Cochran, E. S., Lawrence, J. F., Christensen, C., and Jakka, R. S. (2009). "The quake-catcher network: Citizen science expanding seismic horizons." *Seismological Research Letters*, 80(1), 26-30.

Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., and Kuksa, P. (2011). "Natural language processing (almost) from scratch." *Journal of machine learning research*, 12(Aug), 2493-2537.

Conti, M. (2017). "The incredible inventions of intuitive ai. TED talk."

Corvey, W. J., Vieweg, S., Rood, T., and Palmer, M. "Twitter in Mass Emergency: What NLP Can Contribute." *Proc., Proceedings of the NAACL HLT 2010 Workshop on Computational Linguistics in a World of Social Media*, 23-24.

De Iuliis, M., Kammouh, O., Cimellaro, G. P., and Tesfamariam, S. (2019). "Downtime estimation of building structures using fuzzy logic." *International Journal of Disaster Risk Reduction*, 34, 196-208.

Imran, M., Castillo, C., Lucas, J., Meier, P., and Rogstadius, J. "Coordinating human and machine intelligence to classify microblog communications in crises." *Proc., ISCRAM*.

Kammouh, O., and Cimellaro, G. (2018). "Cyber Threat on Critical Infrastructure: A Growing Concern for Decision Makers." *Routledge Handbook of Sustainable and Resilient Infrastructure*, P. Gardoni, ed., Routledge, London, 359-374.

Kammouh, O., Cimellaro, G. P., and Mahin, S. A. (2018). "Downtime estimation and analysis of lifelines after an earthquake." *Engineering Structures*, 173, 393-403.

Kammouh, O., Dervishaj, G., and Cimellaro, G. P. (2017). "A New Resilience Rating System for Countries and States." *Procedia Engineering*, 198(Supplement C), 985-998.

Kammouh, O., Dervishaj, G., and Cimellaro, G. P. (2018). "Quantitative Framework to Assess Resilience and Risk at the Country Level." *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 4(1), 04017033.

Kammouh, O., Noori, A. Z., Cimellaro, G. P., and Mahin, S. A. (2019). "Resilience Assessment of Urban Communities." *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 5(1), 04019002.

Kammouh, O., Noori, A. Z., Taurino, V., Mahin, S. A., and Cimellaro, G. P. (2018). "Deterministic and fuzzy-based methods to evaluate community resilience." *Earthquake Engineering and Engineering Vibration*, 17(2), 261-275.

Kang, B., and Choo, H. (2016). "A deep-learning-based emergency alert system." *ICT Express*, 2(2), 67-70.

Kazakci, A. O. (2014). "Conceptive artificial intelligence: Insights from design theory." *arXiv preprint arXiv:1404.0640*.

Kilburn, C. R., De Natale, G., and Carlino, S. (2017). "Progressive approach to eruption at Campi Flegrei caldera in southern Italy." *Nature communications*, 8, 15312.

Laaksonen, J., and Oja, E. "Classification with learning k-nearest neighbors." *Proc., Proceedings of International Conference on Neural Networks (ICNN'96)*, IEEE, 1480-1483.

Moon, S.-H., Kim, Y.-H., Lee, Y. H., and Moon, B.-R. (2019). "Application of machine learning to an early warning system for very short-term heavy rainfall." *Journal of Hydrology*, 568, 1042-1054.

Musaev, A., Wang, D., and Pu, C. "LITMUS: Landslide detection by integrating multiple sources." *Proc., ISCRAM*.

Orcutt, M. (2016). "Are face recognition systems accurate? depends on your race." *MIT Technology Review*.

Perez., C. E. (2016). "The alien style of deep learning generative design. www.medium.com."

Radianti, J., Granmo, O.-C., Sarshar, P., Goodwin, M., Dugdale, J., and Gonzalez, J. J. (2015). "A spatio-temporal probabilistic model of hazard-and crowd dynamics for evacuation planning in disasters." *Applied Intelligence*, 42(1), 3-23.

Ramil, A., López, A., Pozo-Antonio, J., and Rivas, T. (2018). "A computer vision system for identification of granite-forming minerals based on RGB data and artificial neural networks." *Measurement*, 117, 90-95.

Shueh, J. (2016). "One concern: Applying artificial intelligence to emergency management. www.govtech.com."

Tresp, V., Bundschus, M., Rettinger, A., and Huang, Y. (2006). "Towards machine learning on the semantic web." *Uncertainty reasoning for the Semantic Web I*, Springer, 282-314.

Tresp, V., Bundschus, M., Rettinger, A., and Huang, Y. "Towards Machine Learning on the Semantic Web." Springer Berlin Heidelberg, 282-314.

Tutun, S., Khasawneh, M. T., and Zhuang, J. (2017). "New framework that uses patterns and relations to understand terrorist behaviors." *Expert Systems with Applications*, 78, 358-375.

Xu, X., Xu, S., Jin, L., and Song, E. (2011). "Characteristic analysis of Otsu threshold and its applications." *Pattern recognition letters*, 32(7), 956-961.

Xu, Z., Zhang, H., Hu, C., Mei, L., Xuan, J., Choo, K. K. R., Sugumaran, V., and Zhu, Y. (2016). "Building knowledge base of urban emergency events based on crowdsourcing of social media." *Concurrency and Computation: Practice and experience*, 28(15), 4038-4052.

Zhang, G., Wang, Z., and Chen, Y. (2018). "Deep learning for seismic lithology prediction." *Geophysical Journal International*, 215(2), 1368-1387.