Optimising road network recovery from major multi-day disruptions

Antonio Pellicer-Pous and Neil S Ferguson,
Department of Civil and Environmental Engineering, University of Strathclyde (Glasgow).

1 Introduction

Road transport network is an essential infrastructure which contributes to the development of a country as it allows the movement of people and goods. Local and regional economy are highly dependent on the transport network as they need to use road infrastructure to provide services to other communities. A recent report prepared by Peeling et al. (2016) that assesses the contribution to the Trunk Road Network (TRN) to the economy in Scotland, demonstrates that Scottish major roads contribute around £1.38 billion to their economy each year through public transport, road freight and road construction and maintenance. The disruption of these roads may result on significant economic and social consequences. In fact, according to a research commissioned by Highlands and Islands Enterprise (HIE) about the transport connectivity and economy of the Scottish areas of Argyll and Bute, road closures produce longer journey times, additional delays and costs for businesses, among other consequences (Ekosgen, 2016).

Road infrastructure is highly exposed to the impact of natural and human-made hazards. Although Scotland is fortunate not to face the same degree of risk from natural hazards compared to other countries, its road network is still vulnerable to the impacts of weather-related hazards. Past events evidence the weakness of the road infrastructure in the presence of extreme weather conditions. As an example, in August 2004, an intense rainfall led to a large number of landslides across Scotland. Some roads were closed for several days, affecting the access to and from remote communities, producing social and economic impacts (Winter *et al.*, 2013, 2016). Climate change models for Scotland predict an increase in precipitation in winter (UK Met Office, 2018), which means that more landslides and flood events are likely to occur in a near future. In situations such as these, infrastructure owners and operators must make decisions about how best to manage and restore the network with limited resources. Hence, they need to be prepared to face these potential future scenarios and recover from them in an efficient and safe way, causing the minimal disruption to users.

To study the restoration of road networks, it is necessary to introduce the general concept of *resilience.* Even though there is a wide diversity of definitions of resilience in the literature, the most accepted is the one provided by Bruneau *et al.* (2003) which defines resilience as the ability to mitigate hazards, absorb the effects of the impacts, recover from them in a way that minimise social disruption and predict and prepare for future stressors. Whilst network robustness (that is, the ability of a network to withstand stress) has been considered in numerous studies, the recovery of networks has captured less attention among researchers despite its significant importance.

A disruption profile and consequent recovery after a disruptive event is shown in Figure 1. Six phases can be identified: (A) Undisrupted state; (B) Degradation of the system performance after a hazard event; (C) Adaptation of drivers to a new disrupted situation; (D) Stable disrupted state, when the network achieves a new equilibrium under disrupted conditions; (E) Recovery process when repairs are carried out; (F) New undisrupted state after all repairs are completed. The priority order of repairs and the allocation of repair resources to damaged places can determine the shape of the recovery graph (from C to F).

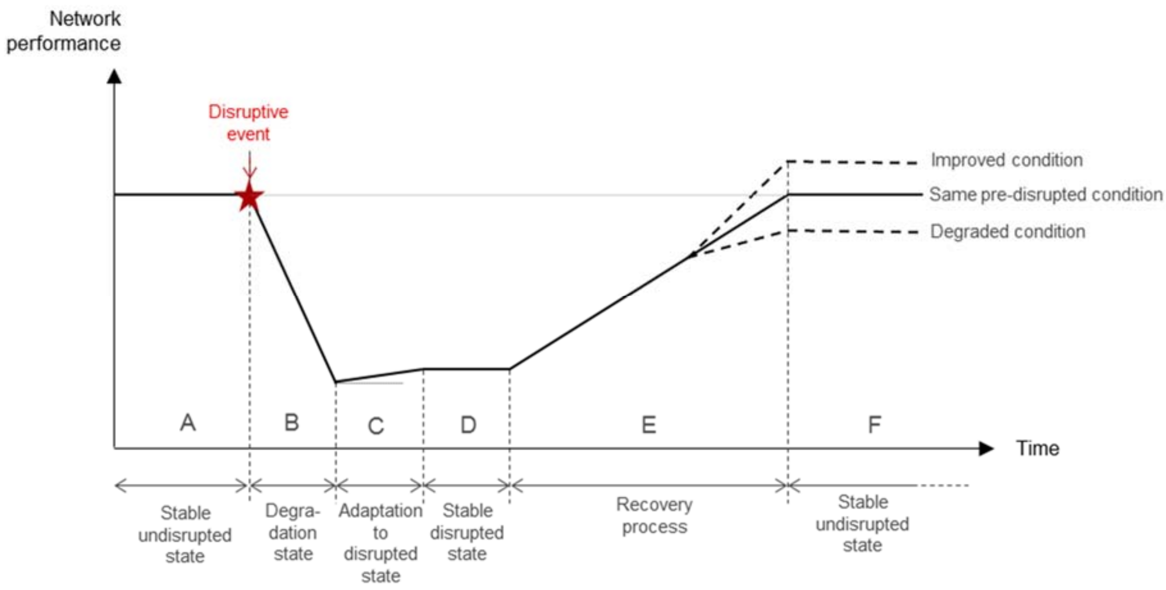

Figure 1: Typical disruption profile and posterior recovery.

Several studies have addressed the problem of obtaining the optimal recovery strategy to certain damaged scenarios. However, the lack of practicality, the low resolution of the model and the assumptions made in each model are unable to capture essential aspects of an actual road network restoration. Also methodologies developed to date are not directly suited to the types of hazard faced by the Scottish road network, which means that further research needs to be done.

In this paper, we present a decision-support tool that can potentially be used by transport authorities and operators to help them make the optimal allocation of resources to damaged roads and prioritise repairs after major disruptions. The novelty of this model lies in the combination of: (1) an infrastructure repair module that simulates repair processes; (2) an improved learning-based traffic module that simulates how drivers make day-to-day and within-day travel decisions; and (3) an optimisation model that is able to find the 'best' repair strategies based on the values of certain performance metrics. The model offers the possibility of testing different repair strategies so that different daily drivers' behavior can be compared. This driver' behavior can also be altered by providing real-time traffic information (used commonly in Intelligent Transportation Systems, ITS) so that drivers can also make more informed day-to-day decisions. This paper also provides a study based on the Scottish road network that illustrates the usefulness and applicability of this model to potential real case scenarios.

The paper is outlined as follows: Section 2 deals with related work. Section 3 describes in detail the methodology used to create the model proposed on this paper. Section 4 applies the model to a real case scenario based on the Scottish road network. Final conclusions and some directions for further work are proposed in Section 5.

2 Background

A systematic literature review was carried out in order to identify among the literature those studies that modelled the recovery of road networks. Results from the review reveal that current methodologies developed to date are still too simplistic to capture the essential aspects of a real road recovery process. Most models are specially designed to deal with disrupted networks caused by the so-called 'low-probability high-consequence' scenarios (commonly known as "disasters"). Although Scotland is fortunate not to face the same degree of risk from natural hazards compared to other countries, its road network is still vulnerable to the impacts of lower-risk weather-related hazards, such as landslides, floods and high winds. In this sense, it makes these models not directly suited to the types of hazard faced by the Scottish road network.

The modelling of traffic demand is overly too simplistic for dynamic, multi-day situations where drivers have to adapt to unforeseen changes on network conditions. The majority of recovery models uses the assumption of User Equilibrium (UE), originally formulated by Wardrop (1952). This condition assumes that drivers are able to choose the route with the minimal travel cost at every time and therefore, no driver can reduce his/her travel cost by switching to other routes. However, this principle hides a strong assumption: drivers need to have ‘perfect knowledge’ of traffic conditions at all time. Some authors (Faturechi and Miller-Hooks, 2014) justify the usage of this UE assumption only during a long-term recovery process so that drivers have plenty of time to find the best route between repairs. However, it is still not really clear to what extend the boundary of ‘long term’ recovery is defined.

To make it clearer to the reader, the graph of Figure 1 is used to explain the implications of the UE assumption. Under UE conditions, drivers have to adjust their travel decisions so that equilibrium is not disturbed. It means that they can only reach stable states (‘A’ and ‘D’ in graph). The transition between these states is therefore omitted due to the assumption of ‘perfect knowledge’. However, in reality, it is impossible to know (especially before travelling) other drivers’ choices and thus, it is more likely to assume that UE is not reached. Through repetitive actions, after undertaking the same trip day after day, drivers may find the route with minimal travel cost. The more information drivers have, the easier it will be for them to find the best route. This ‘imperfect knowledge’ of travel conditions is considered on states ‘B’ and ‘C’ of the graph in Figure 1. The type of hazards that Scotland face may not require a really long-term recovery process. This means that the usage of UE conditions may reflect unrealistic results as the adaptation process of drivers to new network states cannot be negligible.

As observed, there is still research potential in this field despite the progress made so far. In order to avoid UE-based models, the following Section 2.1 briefly reviews previous models that simulate the process of how humans make day-to-day decisions and learn from them. The idea behind these models can be used to model how drivers make day-to-day travel decisions and thus, overcome the limitations that UE-based models have.

2.1 Previous stochastic learning models

One of the first classical learning models was proposed by Bush and Mosteller (1951), using some considerations of Estes (1950). It is based on the idea of representing human behaviour as a process in which an action is chosen between a set of different behavioural alternatives, each one with an associated probability of being selected. When a person needs to make a decision, he/she selects stochastically between these set of alternatives. The ‘reinforcement learning process’ comes into play when the consequences of today’s decision influence future ones (see Figure 2). If the action leads to positive outcomes, then the probability that this action is taken subsequently is increased. If the action leads to a negative outcome, then its probability is decreased. In the psychology literature, this is known as *stochastic learning* (Niemark and Estes, 1967) and in the machine-learning literature, *learning automaton* (Narendra and Thathachar, 1974).

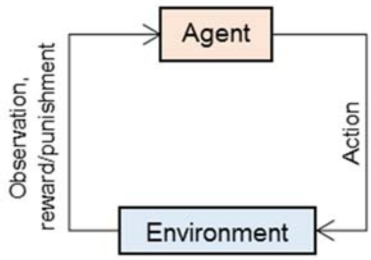

Figure 2: Basic diagram of a Reinforcement Learning system
[Adapted from Sutton and Barto (2017)]

Recently, this theory has been applied to the travel behaviour research area. One of the recent models that applied the stochastic learning automata theory to recreate day-to-day drivers' decisions was the one developed by Ozbay, Datta and Kachroo (2001). The authors implemented a route choice model in which drivers selected their routes based on their travel time from previous days. A year later, the authors modified the model in order to incorporate departure time choice in addition to route choice (Ozbay, Datta and Kachroo, 2002). Wei, Ma and Jia (2014) developed a route choice model following the same principle as the Bush-Mosteller model but providing a more detailed calculation of the expected and experienced travel time. The model developed by de Oliveira Ramos and Grunitzki (2015) added the additional mechanism of updating the drivers' set of routes, so that drivers could learn quickly other faster routes.

However, traditional traffic behavioural models need to be modified in order to be used in the context of dynamic transport disruptions and ITS. Existing stochastic learning approaches have serious limitations in evaluating the interaction between user decisions and external information that drivers may receive during disrupted traffic conditions. The model proposed in this paper tries to overcome these limitations and study the evolution of a traffic system in response to major multi-day disruptions.

3 Methodology

3.1 Modelling framework

This paper presents a framework that provides a modelling tool to assess the effect of different road recovery strategies on traffic behaviour. The model is capable of obtaining the optimal strategy to repair damaged infrastructure, maximising connectivity and minimising overall disruption. The proposed approach integrates a hazard-capacity loss assessment, an infrastructure-based repair process, a stochastic learning traffic assignment model and an optimisation module. The model addresses many limitations of existing recovery models, especially those related to the traffic behaviour modelling, by using recent methodologies in the area of 'artificial intelligence'.

The developed framework is summarised in Figure 3. The occurrence of a hazardous event, such as landslides or floods as the most common weather-related event faced by Scotland, may cause physical damage to certain road sections, resulting on totally or partially road closures and hence, road capacity losses. Depending on the selected recovery strategy, this capacity loss is restored at different speeds and as a result, drivers adapt differently to these distinct recovery strategies. In this sense, the effectiveness of each recovery strategy in terms of connectivity and traffic performance can be assessed. The optimal recovery strategy will be that one that maximises or minimises the proposed metrics. Each of these modules are explained more in detail in the following sections. Despite the importance of assessing the functional loss of a network after a hazardous event, this sub-module 'hazard-damage to infrastructure' is not considered in this paper. Further work is planned to be done in this area so that the proposed approach can be totally completed.

This integrated method, implemented in MATLAB (The MathWorks Inc., 2018), provides an advanced decision-support tool that helps transport planners and authorities to analyse and select the optimal recovery strategy between a broad range of potential recovery solutions. This model is especially designed to deal with scenarios in which there exist multiple damages across the road network and limited repair resources. In these cases, transport planners need to decide where to optimally allocate their limited resources across the network.

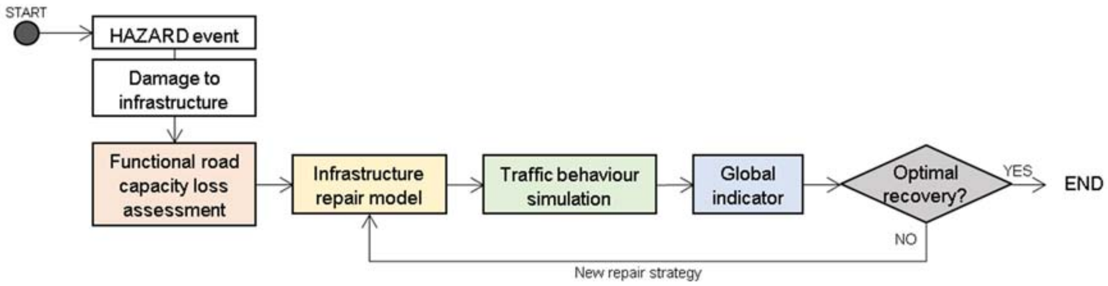

Figure 3: Methodological framework of the proposed model.

3.2 Infrastructure repair model

The process of repairing road infrastructure in real life is formed by a complex sequence of actions that requires interactions between different stakeholders. It is almost impossible to capture all aspects of a real repair process. However, the model proposed in this section tries to identify those essential aspects of a recovery process and implement them on a repair algorithm.

The restoration process of each damaged road proposed on this paper is modelled following a three-way relationship *damage-capacity-time,* shown in Figure 4. In this model, day-to-day road disruptions are represented as a drop of road capacities. Repairs, which are assumed to be done overnight, are carried out in order to restore road capacities to their pre-disrupted levels. Different repair strategies will provide different evolution of capacity-disruption levels over time.

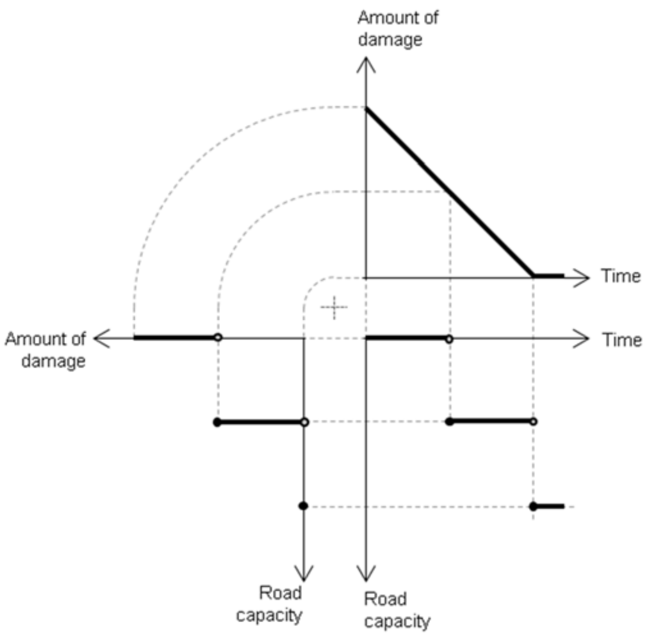

Figure 4: Damage-Capacity-Time relation proposed in this model.

Damage-Time relationship. Damages to road sections are assumed to be known from post-event inspections and can be quantified in units of 'restoration work'. These damaged roads require some physical resources in limited supply (resource-constraint approach) to be repaired. Repairs on each road are represented as a single overall task with an associated amount of restoration work to be done. This work is carried out by one or more repair teams which are formed by a certain number of personnel, plant and equipment. Each repair team has an overall work pace which is measured in terms of 'work/day'. This model assumes a linear relationship between the number of repair teams and the restoration efficiency as shown in Figure 5. This means that more than one repair team attending the same incident can speed up the recovery process until a certain point which is when a saturation of repair teams is reached. In this sense, depending on how many resources are assigned to each damaged road, the recovery can be linearly faster or slower. Resources are assigned to damaged places following the priority order of repairs. If spare resources are available, simultaneous repairs on different places can be carried out.

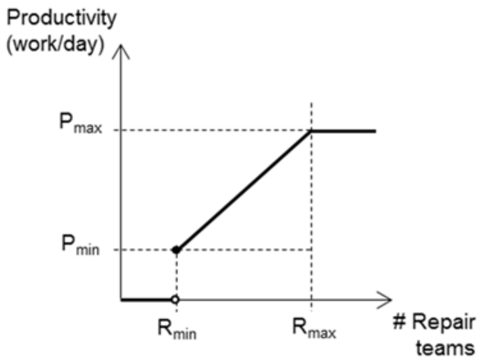

Figure 5: Saturation of repair efficiency

Damage-Capacity relationship. As a result of completion of specified tasks, the capacity of certain roads is augmented. Transport operators are the ones responsible to decide at which point the road can be totally or partially operational again, ensuring the safety of all user at all times. As observed in Figure 4, this model assumes a two step-wise relation in which capacity is restored from 0% to 50%, when some damaged is repaired, and then to 100%, when damage is completely repaired.

Capacity-Time relationship. The output of this repair module is this evolution of road capacities over time (see time-variant capacity graph in Figure 4). It can easily be obtained if the two previous relations (*damage-time* and *damage-capacity*) are known. This *capacity-time* relation is used to simulate the variation of road capacities over time as a result of the considered repairs.

3.3 Stochastic learning traffic model

The aim of this section is to describe a departure time and route choice model that is able to simulate the adaptive process of drivers to day-to-day network changing conditions. The proposed sub-model is inspired by the early work of Ozbay, Datta and Kachroo (2001, 2002) and the recent work of Wei, Ma and Jia (2014) in the modelling of travel behaviour as a form of 'machine learning'. In the field of "artificial intelligence", this model adopts a reinforcement learning (RL) algorithm that allows agents (drivers) to take actions (travel options) and interact with the environment (traffic conditions), so that good actions are rewarded and bad ones are punished. This algorithm is especially suited to deal with uncertain environments which makes it appropriate to be applied to travel behaviour modelling.

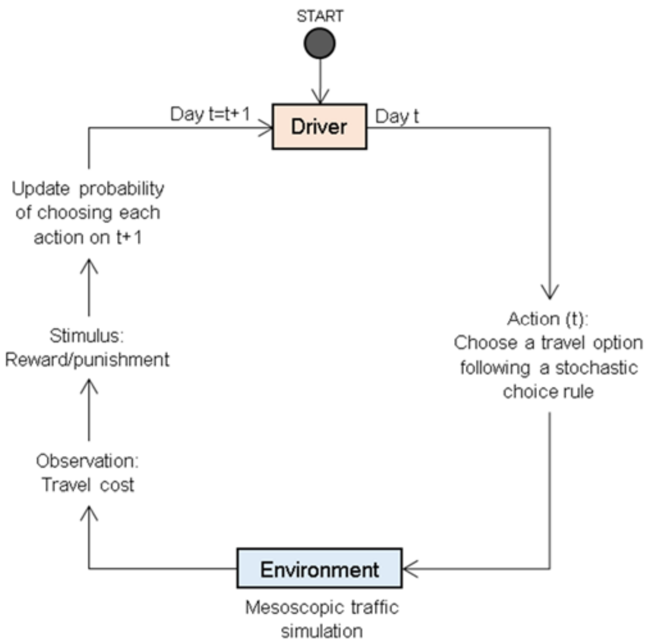

Figure 6: Process of drivers' decision making.
(See the analogy with the RL algorithm on Figure 2)

The drivers' decision making process considered on this model is shown on Figure 6. Each day, drivers select stochastically their travel option (simultaneous departure time and route choice). The movements of vehicles on the network is then simulated using a mesoscopic traffic simulator. The advantage of this model compared to other similar models lies on the incorporation of on-board re-routing dynamics. If drivers face a disruption, they can either use any alternative route used in their past travel experience or follow the diversion route set in place if the road is shut. Drivers can also decide to abandon the trip and return home if the journey is expected to take a long time.

The possibility of providing external information to some drivers is also included in this model. Advanced Traveller Information Systems (ATIS), as part of Intelligent Transport Systems (ITS), have provided extra travel information via websites, GPS devices, mobile phones, Variable Message Signs VMS, etc. With this information, drivers are able to adjust their trip departure time and route choices to real-time traffic conditions. The different types of information considered in this model and their potential behavioural response are included in Table 1.

Table 1: Different types of information and its potential for influencing travel behaviour

	Type of information			
Behavioural response	Pre-trip	VMS	GPS (trip planning)	GPS (on-route)
Avoid totally/partially closed roads	✓	✓	✓	✓
Change route (before travelling)			✓	✓
Change route (on-board)		✓		✓
Change departure time	✓*			
Trip suppression	✓	✓	✓	✓

* This is not included yet in the current model, but it is planned to be added in a new version.

The reinforcement learning algorithm allows drivers to improve their travel decision from day to day and learn from their mistakes. This is calculated based on a 'stimulus' function. It compares the expected travel cost (ETC) of the selected option and the experienced travel cost (PTC) after travelling. If the PTC is less than the ETC, then this driver is more likely to repeat the same option for the next day.

3.4 Evaluation of repair alternatives – Functionality based metrics

Resilience is quantified based on the evolution of the functionality of transport networks over time. Different ways of calculating resilience values are found in the literature (Wang and Yodo, 2016; Sun, Bocchini and Davison, 2018; Vishnu, Kameshwar and Padgett, 2019). In this paper, resilience value is defined as the area under the network functionality curve, as shown in Figure 7. This resilience value is used as an indicator of the effectiveness of each recovery strategy which is applied to the damaged network.

Two functionality metrics have been used in this model to monitor the performance of the system: (1) Connectivity, defined as a measure of the extent to which the nodes of a network are connected to one another; (2) Travel expenditure, that quantifies the average travel time on each day.

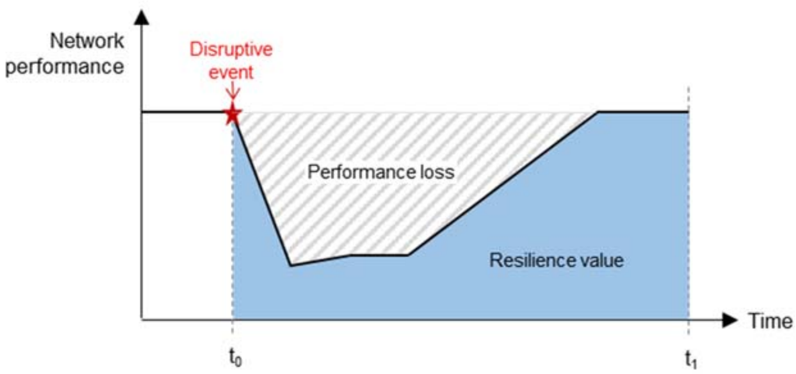

Figure 7: Graphical representation of the resilience value calculation.

3.5 Optimisation module

A discrete bi-objective optimisation problem is proposed in Eq. 1. It aims to find the ‘best’ repair sequence and allocation of resources among all feasible solutions, with the limitation of available repair resources, maximising network connectivity and minimising overall disruption.

$$f(x) = \begin{cases} \min & R_{travel}(x) \\ \max & R_{conn}(x) \end{cases} \quad \textbf{(1)}$$

Subject to:
Limitation of resources
x, integer values

where,
$R_{travel}(x)$ and $R_{conn}(x)$, are the resilience values associated to the average travel time metric and connectivity metric, respectively.
x, decision variables

Various types of heuristics algorithms (e.g. genetic algorithms, simulated annealing algorithms, swarm optimization, local search algorithms, etc.) have been applied in the literature to solve this type of integer sequential scheduling problems. However, the most commonly used technique to solve complex scheduling problems is the Genetic Algorithms (GA) (Chang, 2018), which is the algorithm proposed for this model.

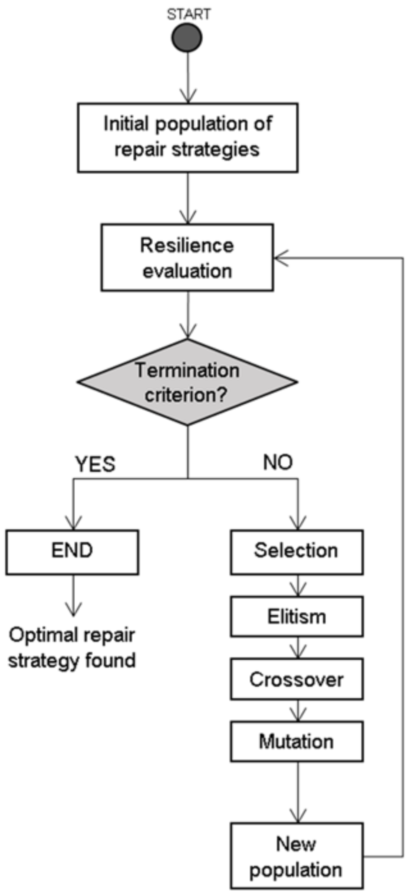

Figure 8: Flowchart of the genetic algorithm.

Genetic Algorithms (GA) are used as a technique to solve the optimisation model. Basically, it simulates the Darwinian principle of evolution in order to find better solutions. The process of a typical GA is shown in Figure 8. Initially the GA operates on a population of *R* repair strategies. For each possible solution, the repair-traffic model is run in order to get a resilience value which indicates how good the solution is to this particular problem. After all initial population is evaluated, a process of fitness-based selection, elitism, crossover and mutation is carried out in order to generate a new population of repair strategies. This is an iterated process in which the population evolves through successive generations until a stopping criterion is reached. A set of 'best' optimal repair strategies are found at the end of the GA process.

4 Case study modelling

To demonstrate the applicability of this methodology, we apply our model to the Scottish road network. Due to its susceptibility to landslide activity (Postance *et al.*, 2017) and its sparse network, the area of study is located on the Highland region (Scotland). The area is divided into 24 internal zones and 4 external zones, which represent the rest of Scotland. Figure 9 shows the network configuration and the study area.

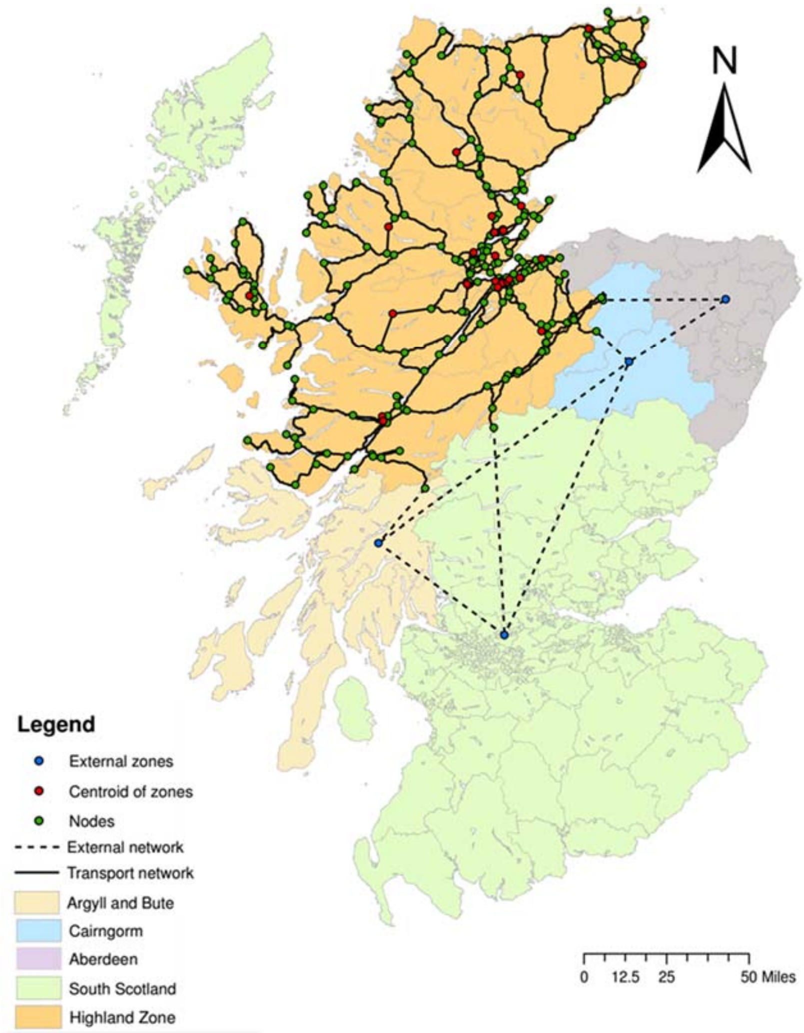

Figure 9: Study area and road network (Scotland).

The road network is modelled as a graph where road intersections are represented by 260 nodes and roads by 728 links. Network topology and attributes are obtained from Ordnance Survey (OS) data. Only roads classified as strategic, primary and secondary are considered in the study. To speed up the simulation process, vehicles are grouped into packets of 20. The Origin-Destination (OD) trip matrix for the morning peak period (7am to 10am) is obtained from the National Trip End Model (NTEM) data, via TEMPro software (Department for Transport, 2017). Each trip departs/arrives from/to the centroid of each zone, which represents the 'centre of activity'. Trips are loaded to the transport network through connectors joining centroids to some nearest nodes. A total of 396 OD pairs are included in this model. It is assumed that 25% of drivers are using on-board GPS navigation, other 25% are using pre-trip information to obtain the shortest path prior to departure and the rest of drivers have no external information.

The model is performed on a computer with 8 GB memory and a quad-core 3.3 GHz Intel i5-3550 processor. At the time of writing this paper, the computational time required to run each simulation is too high. This makes the GA optimisation unfeasible for a population of 100 repair strategies and several generations. It is planned to speed up the process by either using a High Performance Computing, improving the algorithm or using Parallel Computing. However, to overcome this limitation and show the applicability of this model to real scenarios, only three random repair strategies are chosen. A comparison between these strategies will be made and future work will include results from the optimisation module. Section 4.1 provides a more detailed explanation of the damage scenario and proposed repair strategies.

4.1 Disrupted scenario and repair strategies

The case study assumes 6 incidents across the network (see Figure 10). These roads are damaged on day 8 and need to be completely shut until certain damage is repaired. Only 20 repair teams are available to repair these incidents. The proposed repair strategies which include the priority order of repairs and the maximum number of teams that can attend each damaged place are included in Table 2.

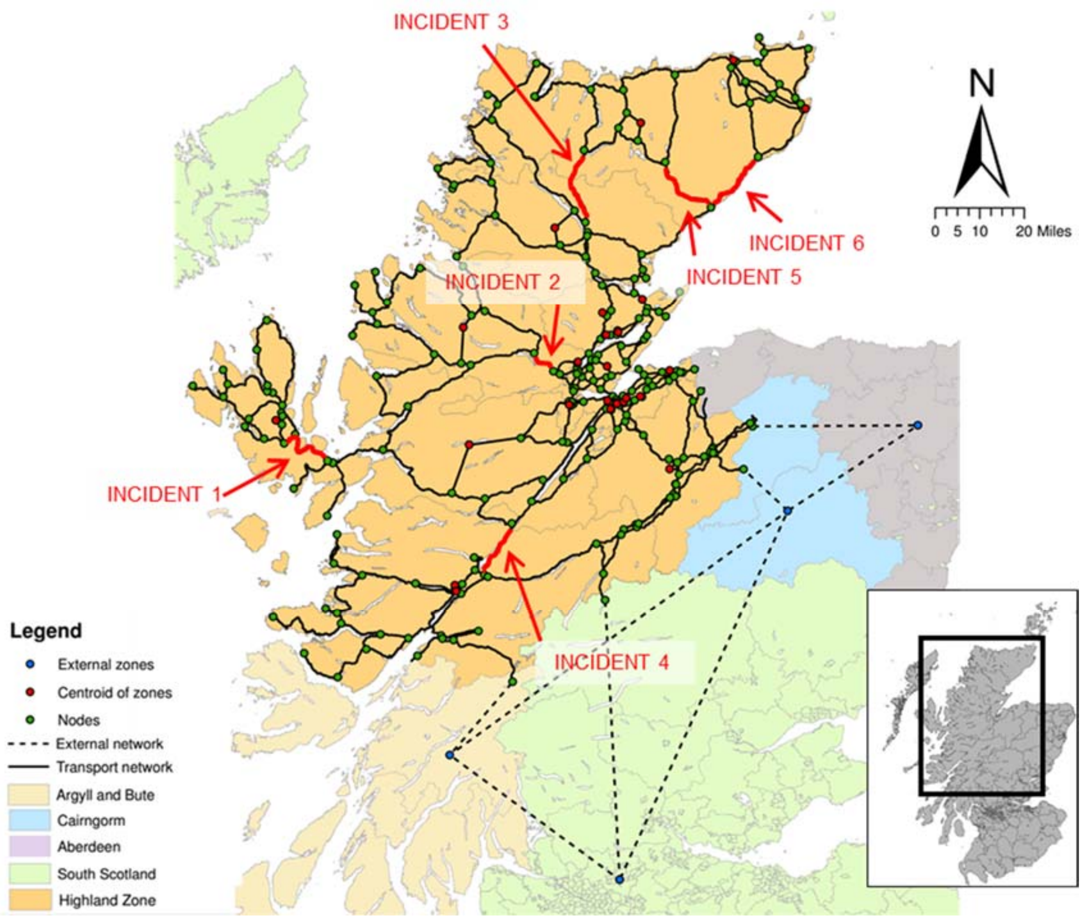

Figure 10: Location of damaged roads

Table 2: Damage values and proposed repair strategies

		Repair strategy 1		Repair strategy 2		Repair strategy 3	
	Damage (value)	Priority order	# rep. teams	Priority order	# rep. teams	Priority order	# rep. teams
Incident 1	Extreme (40)	1	10	6	8	3	6
Incident 2	Slight (15)	2	4	3	7	5	3
Incident 3	Moderate (20)	3	4	5	4	1	10
Incident 4	Slight (15)	5	5	4	7	2	10
Incident 5	Moderate (30)	4	6	1	4	6	8
Incident 6	Slight (15)	6	7	2	5	4	5

4.2 Results and discussion

Results from the evaluation of the proposed repair strategies are included in this section. The evolution of the network performance over time is shown in Figure 11. Due to high computational costs, only one simulation is represented for each repair strategy. Note that, due the stochastic nature of the problem, the model needs to be run more than once to get more representative results.

The graph (11.a) of Figure 11 represents the evolution of travel time over the simulation period. During the first 7 days, there is a warm-up phase in which drivers are getting information (creating on their own experience) about daily traffic conditions and improving their departure time and route choices through the learning process. Only when the overall performance is stable (approx. day 6 or 7), it is assumed that drivers have found a preferred travel option. Note that this does not mean that the network is under Wardrop's user equilibrium. Some of them may have chosen an option that, in reality, could not have been the best among the rest of options. This evidences the 'imperfect knowledge' that drivers have if no external information is provided. On day 8, the closure of some roads produces an increase on the average travel time as drivers have to adapt to new network conditions and take longer routes to reach their destinations. Even after all roads are repaired, drivers are still adapting to the new environment and choosing their most appropriate travel option for each day. Graph (11.b) and (11.c) are related as the lack of connectivity means that some trips cannot be completed as some parts of the network are isolated. However, graph (11.c) also includes those drivers who decide not to travel by car, take another mode of transport or even abandon their trip in the middle of their journey.

Resilience values are obtained as described in Section 3.4 and these are shown in Figure 12. Each dot represents the resilience value associated to each repair strategy after evaluating the performance of the network. Note that, with the shortly incorporation of the GA optimisation module, more than 100 solutions will be added to this graph, obtaining a Pareto front of possible conflicting repair solutions. As the model tries to minimise the average travel time and maximise the connectivity of the network, Repair Strategy 2 can already be dismissed because it is dominated by Repair Strategy 3 which has lower travel time and higher connectivity value. The decision has to be made between Repair Strategy 1, which has higher connectivity value, or Repair Strategy 3, which has lower value of average travel time. At this point, transport authorities need to decide between: (a) keeping all network connected as soon as possible even if it requires to increase slightly the overall travel cost of all drivers, or (b) give more importance to the overall travel cost even if a few drivers cannot reach their destinations during a few days.

The model also provides a Gantt chart diagram with the assignment of repair resources for each repair strategy. The corresponding chart for this particular case study is included in Figure 13.

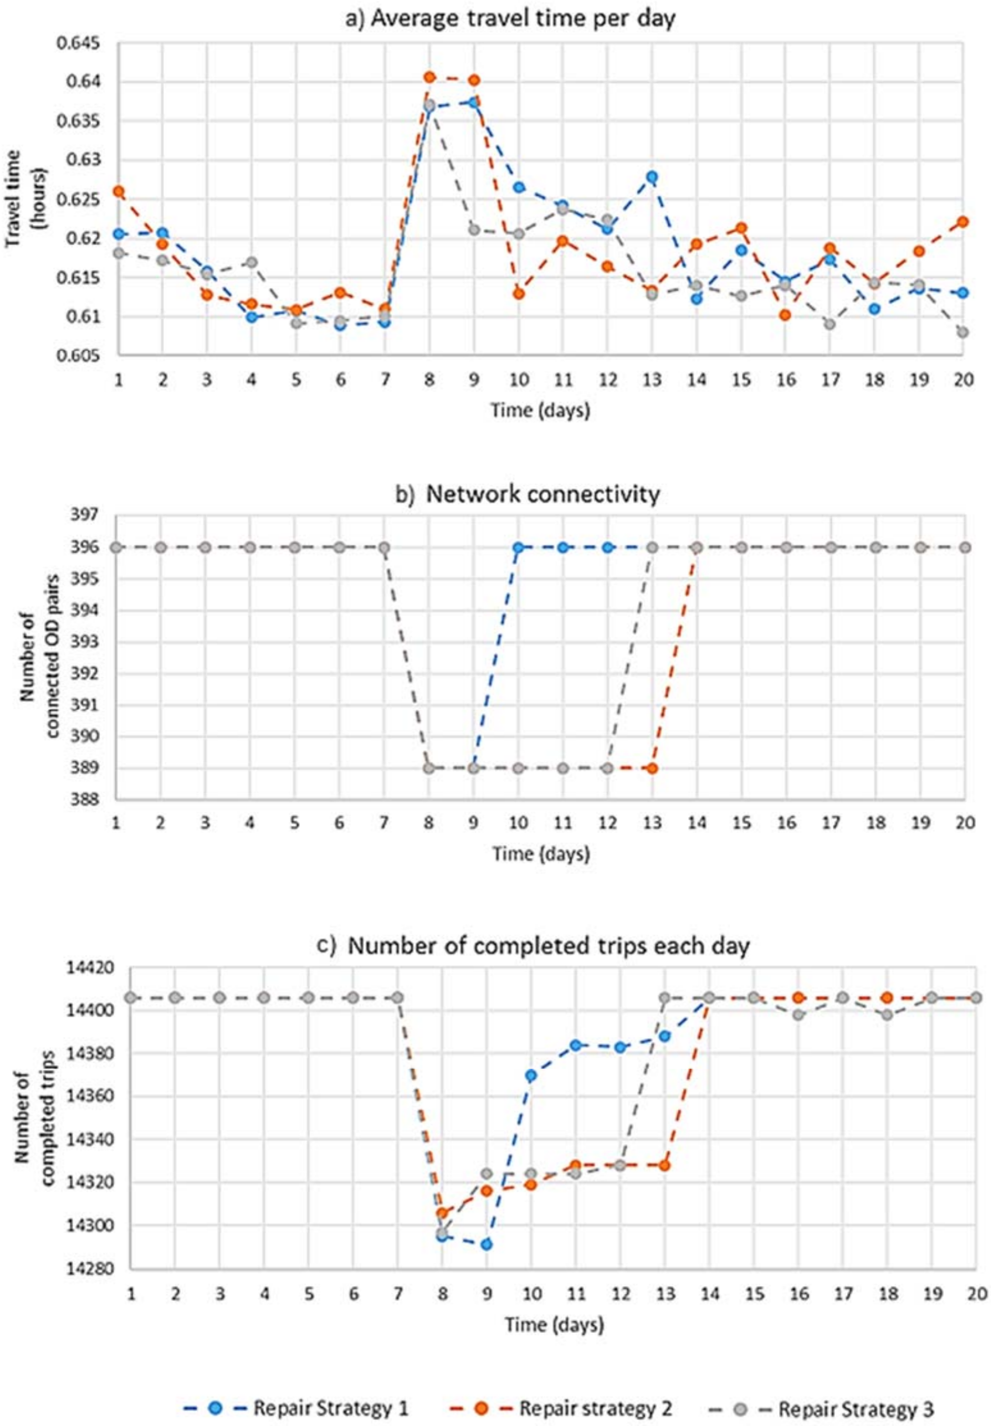

Figure 11: Network performance over time measured as:
(a) Average travel time; (b) Number of connected OD pairs; (c) Number of completed trips.

Figure 12: Resilience values associated to each objective for each repair strategy.

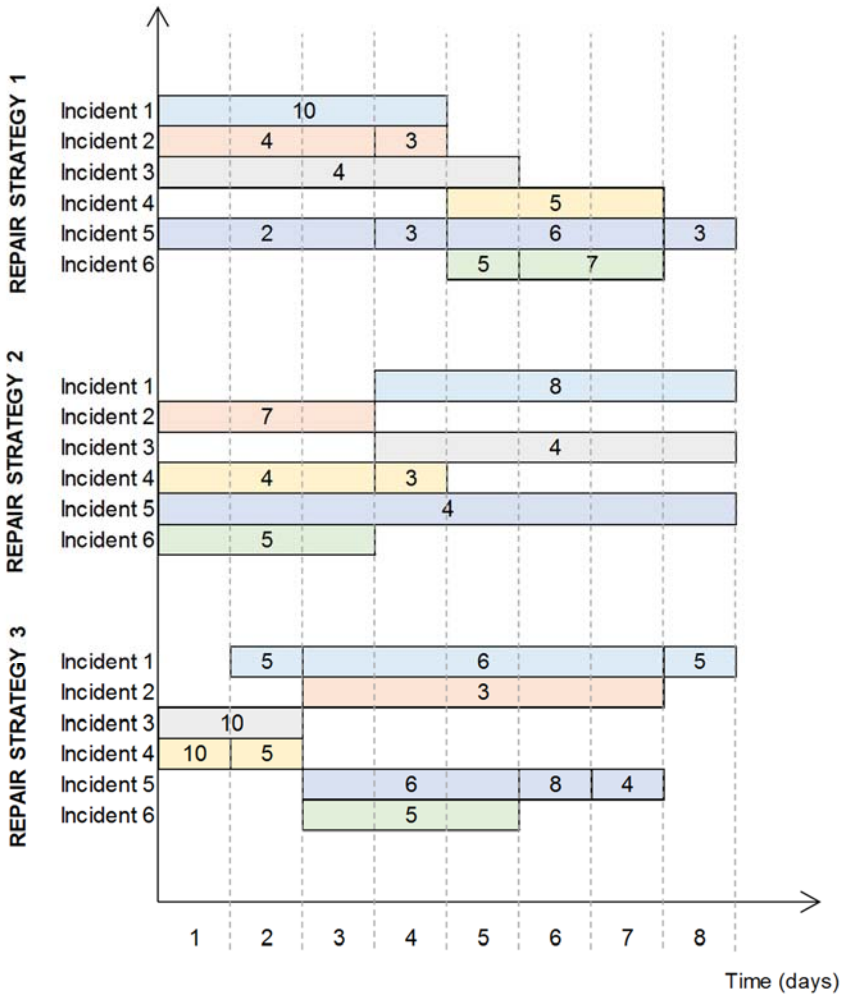

Figure 13: Repair time scheduling and assignment of resources (number inside boxes) for each repair strategy

5 Conclusions

This paper has presented a method to optimise the recovery of road transport networks after major day-to-day disruptions. The model incorporates a module that simulates the process of infrastructure repairs and a stochastic traffic assignment module that models the decision making process of drivers after network changing conditions. The model aims to obtain the optimal recovery strategies that maximise network connectivity and minimise overall traffic disruption. An efficient metaheuristic approach (genetic algorithms) is proposed to find optimal solutions for the recovery problem. This work expands upon previous efforts by introducing a new approach that models the damage-capacity-time relationship and improves the existing learning traffic assignment models to be applicable to low-risk disrupted scenarios. The model was also applied to a real-case scenario based on the Scottish road network. Results from the case study can clearly highlight the potential applicability of this model to evaluate different recovery strategies and optimise the recovery of road networks after multi-day major disruptions.

Inevitably, this paper has its limitations and future work is required to address these issues. Firstly, regarding the infrastructure repair process, a more realistic repair procedure needs to be implemented incorporating different type of resources with different repair efficiencies, cost of repairs and the possibility of bringing external resources at a higher cost for extraordinary situations. Secondly, regarding the traffic model, certain disrupted scenarios can make drivers change their mode of transport. A multi-modal approach could also be considered. Thirdly, future work also requires the implementation of a module that assesses the impact of hazards on road damage and a method to speed up the overall process so that the optimisation module can be run at a higher speed. Finally, the validation of the traffic assignment model using real traffic data is also required.

Although the implementation of a model for improving the road recovery is not an easy task because it involves the participation of a large number of factors and organisations in the process, the model proposed in this paper has provided the first insight into what a road recovery model is and it can potentially be used to help transport authorities and operators in making more optimised decisions.

6 References

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 Von Winterfeldt, D. (2003) 'A Framework to Quantitatively Assess and Enhance the Seismic Resilience of Communities', *Earthquake Spectra*, 19(4), pp. 733–752. doi: 10.1193/1.1623497.

Bush, R. R. and Mosteller, F. (1951) 'A Mathematical Model for Simple Learning', in *Selected Papers of Frederick Mosteller*. New York, NY: Springer New York, pp. 221–234. doi: 10.1007/978-0-387-44956-2_12.

Chang, Y. (2018) *Heuristic approach to network recovery*. University of Oklahome.

Department for Transport (2017) 'TEMPRO—Trip End Model Presentation Program'. Department for Transport. Available at: https://www.gov.uk/government/publications/tempro-downloads.

Ekosgen (2016) *Argyll and Bute Transport Connectivity and Economy. Research Report*. Available at: http://www.hie.co.uk/regional-information/economic-reports-and-research/archive/argyll-and-bute-transport-connectivity-and-economy-research.html.

Estes, W. (1950) 'Toward a statistical theory of learning', *Psychological review*, pp. 94–107.

Faturechi, R. and Miller-Hooks, E. (2014) 'Travel time resilience of roadway networks under disaster', *Transportation Research Part B: Methodological*. Elsevier Ltd, 70, pp. 47–64. doi: 10.1016/j.trb.2014.08.007.

Narendra, K. S. and Thathachar, M. a. L. (1974) 'Learning Automata - A Survey', *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-4(4), pp. 323–334. doi: 10.1109/TSMC.1974.5408453.

Niemark, E. D. and Estes, W. K. (1967) *Stimulus Sampling Theory*. San Francisco.

de Oliveira Ramos, G. and Grunitzki, R. (2015) 'An Improved Learning Automata Approach for the Route Choice Problem', in *Agent Technology for Intelligent Mobile Services and Smart Societies*. Springer Berlin, pp. 56–67.

Ozbay, K., Datta, A. and Kachroo, P. (2001) 'Modeling Route Choice Behavior with Stochastic Learning Automata', *Transportation Research Record: Journal of the Transportation Research Board*, 1752(1), pp. 38–46. doi: 10.3141/1752-06.

Ozbay, K., Datta, A. and Kachroo, P. (2002) 'Application of stochastic learning automata for modeling departure time and route choice behavior', *Transportation Research Record 1807: Traveler Behavior And Values*, pp. 154–162.

Peeling, J., Palmer, D., Booth, C. and Abell, R. (2016) *The Value of the Trunk Road Network to Society and the Economy in Scotland*. Available at: https://www.transport.gov.scot/publication/the-value-of-the-trunk-road-network-to-society-and-the-economy-in-scotland/.

Postance, B., Hillier, J., Dijkstra, T. and Dixon, N. (2017) 'Extending natural hazard impacts: an assessment of landslide disruptions on a national road transportation network', *Environmental Research Letters*, 12(1). doi: 10.1088/1748-9326/aa5555.

Sun, W., Bocchini, P. and Davison, B. D. (2018) 'Resilience metrics and measurement methods for transportation infrastructure : the state of the art', *Sustainable and Resilient Infrastructure*. Taylor & Francis, 9689, pp. 1–33. doi: 10.1080/23789689.2018.1448663.

Sutton, R. S. and Barto, A. G. (2017) *Reinforcement Learning: An Introduction*. Second Edi. The MIT Press.

The MathWorks Inc. (2018) 'MATLAB-R2018b'. Natick, MA.

UK Met Office (2018) *UKCP18 Headline Findings*. Available at: https://www.metoffice.gov.uk/binaries/content/assets/mohippo/pdf/ukcp18/ukcp18-headline-%0Afindings.pdf .

Vishnu, N., Kameshwar, S. and Padgett, J. E. (2019) 'A framework for resilience assessment of highway transportation networks', in *Routledge Handbook of Sustainable and Resilient Infrastructure*, pp. 216–237.

Wang, P. and Yodo, N. (2016) 'Engineering Resilience Quantification and System Design Implications: A Literature Survey', *Journal of Mechanical Design*, 138(c). doi: 10.1115/1.4034223.

Wardrop, J. G. (1952) 'Some theoretical aspects of road traffic research', in *Proceedings of the institution of civil engineers*, pp. 325–362.

Wei, F., Ma, S. and Jia, N. (2014) 'A Day-to-Day Route Choice Model Based on Reinforcement Learning', *Mathematical Problems in Engineering*, 2014, pp. 1–19. doi: 10.1155/2014/646548.

Winter, M. G., Harrison, M., Macgregor, F. and Shackman, L. (2013) 'Landslide hazard and risk assessment on the Scottish road network', *Geotechnical Engineering*, 166(GE6), pp. 522–539. doi: http://dx.doi.org/10.1680/geng.12.00063.

Winter, M. G., Shearer, B., Palmer, D., Peeling, D. and Sharpe, J. (2016) 'The Economic Impact of Landslides and Floods on the Road Network', *Procedia Engineering*. Elsevier B.V., 143(Ictg), pp. 1425–1434. doi: 10.1016/j.proeng.2016.06.168.