Using machine learning to examine freight network spatial vulnerabilities to disasters: A new take on partial dependency plots

Paul M. Johnson[1*], William Barbour[1], Janey V. Camp[1], and Hiba Baroud[1]

[1]Department of Civil and Environmental Engineering, Vanderbilt University
[*]Corresponding author: Paul M. Johnson; paul.m.johnson@vanderbilt.edu; 2301 Vanderbilt Place PMB 35183 Nashville, TN, USA 37240-0002

Abstract

Analyzing transportation network vulnerabilities to disruptions is crucial for society to maintain commodity flows across the globe. However, most vulnerability analyses focus on impacts that arise from the deterioration of single network components, which can overlook spatial correlations between multiple components that manifest during area-spanning disruptions, such as those stemming from natural hazards. Here, we demonstrate an intuitive approach for inferring spatial vulnerabilities to area-spanning disruptions. In particular, we show how partial dependency plots derived from gradient boosting machines trained on routing simulations can be used to depict the average effect a disruption's location has on impacts while controlling for other input variables and spatial interdependencies embedded in the network. Although we demonstrate our approach for Middle Tennessee's intermodal road and rail freight transportation network, our framework can easily be applied to other networks and regions.

Keywords: Network vulnerability; transportation disruption analysis; partial dependency plot; gradient boosting machine; freight simulation

1 Introduction

1.1 Background

The US multimodal freight transportation network is an extensive collection of road, rail, air, waterway, and pipeline transports that is vital for maintaining commodity flows throughout the domestic and global economy (Darayi, Barker, and Nicholson, 2019). In 2017, this system handled approximately 18 billion tons, $17.5 trillion, of freight (USDOT, 2019). Disruptions to this intermodal network, such as those resulting from extreme weather events, infrastructure failure, or terrorist attacks, pose a significant threat to its operations and the well-being of communities around the world (Freiria et al., 2015; Ivanov, 2020). As such, researchers and policy makers have found it prudent to examine *vulnerabilities* within this network.

The term *vulnerability* spans many disciplines but generally refers to a system's or entity's susceptibility to harm (UNDP, 2011). A common approach to analyzing vulnerabilities in transportation networks is to examine impacts that result from the disruption or impairment of network components (i.e., links and nodes) (Sullivan et al., 2009; Whitman et al., 2017). Impacts can be evaluated using theoretical measures derived from graph-theory, such as loss in connectivity, or empirical measures, such as estimated increases travel times or distances (A. Chen et al., 1999; Kim et al., 2015; Murray and Grubesic, 2007; Whitman et al., 2017). Effective risk mitigation strategies typically entail hardening network components whose disruptions result in the most adverse outcomes. The range of applications for these types of vulnerability analyses is vast, including but not limited to emergency health preparedness and disaster response planning (Asakura, 1999; Baghalian et al., 2013; Freiria et al., 2015; Murray-Tuite and Mahmassani, 2004; Peng et al., 2014; Sohn, 2006).

Network vulnerability analyses can generally be categorized by the types of scenarios being studied (Haghighi et al., 2018; Murray and Grubesic, 2007). First, worst-case scenario planning entails researchers focusing on what set of disruptions most severely impact the network (Sullivan et al., 2009; Gedik et al., 2014; Wang et al., 2016; Xu et al., 2017). For example, Murray-Tuite and Mahmassani (2004) modeled a game where an "evil entity" tries to cause the worst possible situation for a traffic management agency. Second, case-specific scenarios involve researchers investigating impacts from a particular event(s). For example, Cho et al. (2001) and Tatano

and Tsuchiya (2008) examine impacts to urban transportation networks resulting from prescribed earthquake hazards likely to affect their areas of interest. Lastly, full-range scenarios involve researchers scanning all possible link removals in the network and determining their subsequent impacts (Jenelius and Mattsson, 2012). For example, Sohn (2006) studied the effects of systematically degrading every link in Maryland's highway network.

While each approach has its merits, full-range scenarios are particularly useful for gaining a holistic understanding of vulnerabilities in a network (Jenelius and Mattsson, 2012; Sugishita and Yasuo Asakura, 2021). Traditionally, these studies involve researchers enumerating all components of a network and then examining the effects of degrading each network component one at a time (Haghighi et al., 2018; Jenelius and Mattsson, 2012; Snyder and Daskin, 2007; Wang et al., 2016). This single-component approach can be appropriate for modeling localized disruptions, such as bridge failures or car accidents (Whitman et al., 2017). However, it fails to capture spatial correlations between network components that are pertinent for area-spanning disruptions, such as those stemming from extreme weather events (Haghighi et al., 2018; Jenelius and Mattsson, 2012; Calatayud et al., 2017; Sugishita and Yasuo Asakura, 2021).

Consequently, several studies have expressed the need to consider simultaneous disruptions of multiple components when analyzing network vulnerabilities (Jenelius and Mattsson, 2012; Xu et al., 2017; Haghighi et al., 2018). For example, Patterson and Apostolakis (2007) noted that risks of terrorist attacks on the Massachusetts Institute of Technology campus' power and utility networks differed when they examined components on an area-spanning basis versus separate from one another. Similarly, Jenelius and Mattsson (2012) found that impacts from area-spanning disruptions along Sweden's roadway were mostly affected by total levels of inbound and outbound traffic, as opposed to cases of single-link failures where proximate network redundancies and segment flows were the main factors.

However, enumerating and evaluating all possible combinations of multi-component disruptions can be computationally impractical for full-range scenarios, especially given a large and/or complex network (Chow and Regan, 2014; Wang et al., 2016; Xu et al., 2017). One way to alleviate this problem is for researchers to limit the number of scenarios being considered, as Jenelius and Mattsson (2012) did by representing area-spanning scenarios as a finite set of grid boxes. One downside with this type of approach is that some important scenarios may be inadvertently left out of the analysis (Wang

et al., 2016; Xu et al., 2017). Alternatively, Wang et al. (2016) and Xu et al. (2017) developed optimization techniques to identify critical network components during simultaneous, multi-component disruptions without needing to scan all possible combinations of events. The drawback of these optimization techniques is that although the impacts of all scenarios are technically considered, they are only evaluated for the best and/or worst case-scenarios (i.e., they establish a lower and/or upper bound on impacts). Consequently, accurate interpolations of impacts from other scenarios are largely infeasible. Lastly, researchers can train surrogate models on a subset of simulated scenarios in order to predict outcomes of other scenarios at lower computational costs (Chow and Regan, 2014; Hartig et al., 2014). For example, Haghighi et al. (2018) simulated disruptions from several area-spanning earthquakes in Salt Lake County and trained a linear regression on results to interpolate outcomes of other related events. However, this type of approach is limited by the effectiveness of the surrogate model.

1.2 Problem Statement

Analyzing vulnerabilities in transportation networks is crucial for helping societies maintain their livelihoods (Darayi, Barker, and Nicholson, 2019). A proven approach for examining vulnerabilities in transportation networks is to evaluate impacts from disruptions across a full-range of scenarios (i.e., perform a complete network scan). However, for simultaneous, multi-component disruptions, such as those arising from natural disasters, evaluating all possible disruptions is computationally burdensome. There are several different approaches to alleviate this problem, such as limiting the number of scenarios being considered, using sophisticated optimization techniques to obviate the need to specify all scenarios, and employing surrogate models to interpolate simulation results at low computational costs. Each approach has its benefits and drawbacks, and ultimately, "there is no analytical approach of transportation network vulnerability with a systematic consideration and quantification of all possible simultaneous disruptions" (Xu et al., 2017).

In this paper, we certainly do not claim to solve this problem but do demonstrate an intuitive approach for inferring and depicting network vulnerabilities to area-spanning disruptions that helps mitigate some of the drawbacks of prior studies. More specifically, we show how partial dependency plots (PDPs) derived from gradient boosting machine (GBM) surrogate models trained on outputs from a routing simulation can be used to depict the

marginal effect a disruption's location has on impacts while controlling for other input variables and spatial correlations between network components, thereby revealing which areas of the network are most vulnerable to disruptions. The GBMs are a much more sophisticated statistical technique than the linear regressions employed by Haghighi et al. (2018) and are able to accurately predict simulation outcomes and model complex spatial dependencies between network components without making any a-priori, simplifying assumptions. Additionally, the PDPs are able to depict a heat-map of network vulnerabilities at a much higher resolution, both in terms of space and scenarios, than the grid-box approached discussed by Jenelius and Mattsson (2012). Although PDPs have been extensively featured in the machine learning literature, to the best of our knowledge, they have not been used to identify vulnerable components in a network.

2 Methods

Our methodology demonstrates how researchers can use PDPs to depict spatial vulnerabilities in transportation disruption analyses. The PDPs are derived from surrogate models that are trained on the results of a transportation network disruption simulation. Here, we employ GBMs as surrogate models because they are powerful, non-parametric machine learning procedures capable of learning spatial interdependencies in networks and accurately predicting outcomes of complex simulations. However, researchers can use other machine learning methods if they are better suited to their particular analysis. Related, we demonstrate our approach as a proof of concept with Middle Tennessee's intermodal road and rail freight transportation network, but our overall framework can be easily extended to other applications.

The remainder of the methodology section is outlined as follows. Section 2.1 describes how we constructed Tennessee's intermodal freight network. Section 2.2 discusses how we simulate disruption scenarios on this network. Section 2.3 shows how we trained GBM surrogate models on the results of the disruption simulations. Section 2.4 discusses how we use PDPs derived from the GBMs to illustrate spatial vulnerabilities in the network to these disruptions. As mentioned, researchers can explore alternative means of simulating impacts from disruptions and training surrogate models (Sections 2.1 - 2.3) and still find utility in the PDPs.

2.1 Middle Tennessee Intermodal Freight Network

In this section, we give an overview of how we construct the road and rail freight intermodal network for the Middle Tennessee region and discuss how we instantiate routes in our simulation. Exact details for reproducing our analysis can be found in Appendix I.

The road portion of the network is primarily based on data from the Freight Analysis Framework (FAF) (Ford, 2017). With the FAF data, we are able to create a georeferenced road network that includes edge characteristics such as expected traversal times and corresponding flow metrics such as freight volume. The road network is depicted in Figure 1 (green lines).

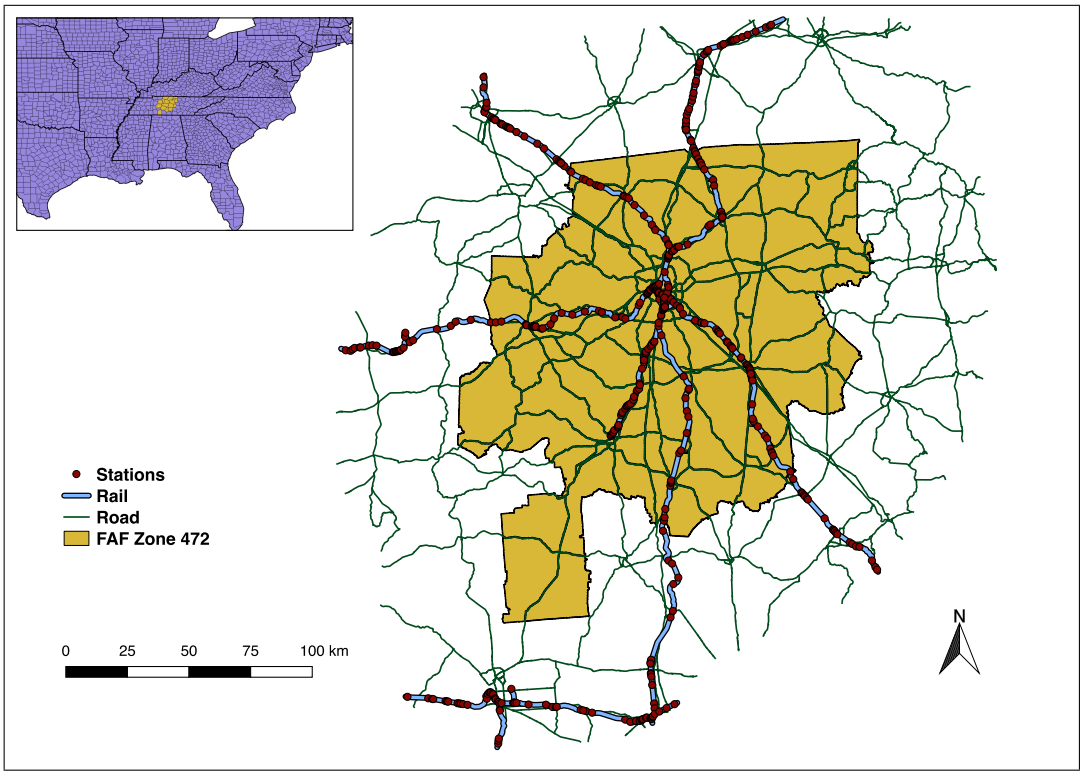


Figure 1: Middle Tennessee Road and Rail Network

We instantiate road freight routes in the simulation at various parts of the network. For major interstates, we use the freight volume reported on each interstate leading into the region, separated by direction, as the originating volume. These originating flows are then distributed to a set of destinations,

comprised of each interstate leading out of the region (through-freight) and a random sample of highway points within the region (terminating routes), based on reported Nashville freight data (Organization, 2010). In total, there are 60 unique interstate-origin flows. Non-interstate origins are determined by randomly selecting 100 non-interstate highway segments in the region. For each of these origins, flows consisting of total reported freight volume on that segment are uniformly distributed to a random selection of ten other non-interstate highway segments as destinations, which results in 1000 distinct OD flows. These respective allocations create an extensive, empirically based, and steady baseline flow along both the interstate and non-interstate highways without the simulation being overly burdensome from a computational standpoint.

The rail portion of the network is primarily based on private Waybill data, which contain records of carloads, tons shipped, commodity identifies, and intermodal capabilities of rail cars that pass through the Middle Tennessee region. However, since each Waybill entry corresponds to an individual Origin-Destination (OD) shipment, we aggregate records by unique legs of the rail network to get segment-level flow estimates in order to match the granularity of the road flows. The US Department of Transportation (USDOT) has recognized the necessity of this task and developed a routing tool to specifically work with the Waybill data to accomplish this task (Wright and Baker, 2017). We use this tool to map rail freight flows to corresponding segments defined by the USDOT, shown in Figure 1 (blue lines).

We instantiate the rail freight flows for the simulation by making the reasonable assumption that they originate on the network edges at the boundary of the region and at major rail yards within the study area. Destinations are simply all other origin points, resulting in 42 unique rail flows. The total originating volume at each point is taken from the Waybill-reported volume. This volume is split between the destination points proportionally, according to their relative originating volumes.

It should be noted that instead of using segment-level commodity flows to ground both the road and rail portions of the network, we could have elected to use the more granular, OD-level Waybill records to ground the rail portion. However, doing so would complicate the routing logic without adding additional fidelity to the road portion of the simulation. Ideally, we would have OD-level data for both rail and road networks on which to base our simulation.

Using ArcGIS' network analysis, we then merge the road and rail networks

into a single, intermodal network, with rail stations serving as the topological nodes that link the two modes of transport (red dots, Figure 1). At these intermodal switching points, we assume the time cost of transfer to be one hour based on industry insights. Next, we convert volumes of road and rail commodity flows into a common unit for analysis, ton-miles per hour, to allow for fair comparisons and transitions between modes of transport.

We base edge capacities on this same unit. For road segment capacity, FAF specifies capacities in ton-miles per hour. For rail, we assume capacity to be the current volume reported in Waybill because many of the nation's rail corridors are operating near capacity in terms of trains per day (Systematics, 2007; Dingler et al., 2010).

All road freight traffic is assumed to be intermodal-capable (i.e., all freight in trucks can be transferred to rail cars if needed); however, this situation rarely occurs because road networks have many alternative routes that are easier to traverse than switching modes of transport. Conversely, the proportion of rail traffic that is intermodal-capable is calculated as the average amount of existing intermodal-capable volume in the region, divided by the total amount of rail volume.

2.2 Simulations

Next, we simulate 24,000 different hazard scenarios and determine their impacts on freight travel times and travel distances. In this section, we first describe the general simulation procedures and then discuss how disruptions cause changes in routing conditions.

2.2.1 Simulation Setup and Baseline Condition

The 1102 total OD flows (1060 road and 42 rail) established in Section 2.1 itemize the amount of freight flowing between various points on the network over a time period of one hour, so dispatching all this volume as if it were a single shipment is not reasonable. As such, we divide the OD freight flows into *blocks* of up to 250 tons per road routes and 100 tons per rail routes, which are distributed uniformly over ten-minute intervals.

At each interval, routes for each flow block are determined simultaneously and non-cooperatively with each other, using Dijkstra's shortest path algorithm for travel times given the current conditions of the network. This routing behavior emulates the concept of traffic-aware routing: preferring

lower-traffic routes with marginally longer distance due to overall time savings (Fotakis et al., 2002). However, routes cannot change once when are dispatched. Formally, this approach is known as Incremental Traffic Assignment, where fixed proportions of travel demand are incrementally assigned at set intervals (Saw et al., 2014).

A network segment's baseline traversal time is determined by the road and rail network data previously discussed, and this traversal time remains constant until volume along a segment exceeds its capacity. When a road or rail network segment exceeds its defined capacity, we penalize its traversal time in proportion to how much the segment is over capacity (i.e., 50% over capacity translates to 50% added travel time). This traversal time function is specified in equation (1) below:

$$T_l = \begin{cases} \gamma \tau_l, & V_l \leq C_l, \\ \gamma \tau_l \dfrac{V_l}{C_l}, & V_l > C_l \end{cases} \tag{1}$$

T_l is the travel time of the network link l based on its current volume V_l, γ is the hazard severity value (discussed later), τ_l is the free flow travel time of link l, and C_l is the capacity of link l. Here, capacity refers to the point at which free flow traffic conditions are degraded, not a hard limit for the maximum volume that can traverse a segment. For the purposes of routing logic in this simulation, imposing a max capacity restriction has the potential to create unresolvable gridlocks by eliminating certain routes altogether. Instead, we elect to limit volume per segment by prohibitively large traversal time penalties.

As such, early simulation flow blocks are routed according to their pure shortest travel time paths because volume on network segments has not reached capacity. Once volume builds up on segments to the point of reaching over-capacity, subsequent flow blocks may choose alternative routes if they take less time to traverse. Ultimately, the simulation runs for a total of 360 time steps, the equivalent of 60 simulation hours, to provide ample time for baseline conditions to reach steady flows before being subjected to disruptions (discussed later).

It should be noted that there are other traffic simulation methods that would be viable alternatives to the approaches we have taken and potentially result in higher fidelity simulations. Traffic assignments could be based on user-equilibrium or stochastic approaches (Saw et al., 2014). Similarly, imposing hard capacity limits on links could be implemented in lieu of time

penalties. Additionally, alternative travel time functions could be used; we used a piece-wise function in order to better represent railway travel times that can tolerate additional traffic without experiencing effects on travel times and because we observed unrealistic effects when using a polynomial travel time model due to low connectivity in some portions of the network. Ultimately, these are all modeling choices that can be further explored but still work with our overall framework.

2.2.2 Disruption Conditions

Hazards are represented as circles with varying size (radius), severity (relative impairment cost), and centroid location (latitude and longitude). Any part of the network that the hazard circle overlaps incurs an impairment cost, modeled as a relative time penalty to traverse that part of the network. This method is similar to the grid-box approach presented by Jenelius and Mattsson (2012). However, the randomized circular hazards allow us to vary the size of the threats and give more precise coordinates to use as inputs to the machine learning models.

In the presence of a disruption, the travel time of affected network segments, whether they are over-capacity or not, is increased proportionally by the severity of the disruption (see, equation (1)). The severity value is randomly chosen from values $\gamma = \{2, 4, 8, 16, 32\}$, and we use this power-law progression to help demonstrate the scalability of the statistical models (i.e., to show that the GBMs can accurately predict "minor" and "severe" disruptions). We apply the severity uniformly across the hazard. This is a modeling choice done to more easily depict where the direct impacts are observed and simply allow indirect effects to propagate outwards from the impact area. Alternatively, one could also model severity as a decreasing function of distance from the centroid of the hazard and/or explore other possibilities. The increased costs of affected network segments (according to equation (1)) are used during origin-destination route computation by Dijkstra's algorithm.

Disruptions are modeled to occur for 20 simulation hours (one third of the overall simulation), beginning 20 hours after the start of the simulation. This delayed onset of the disruption allows simulation flows to stabilize before effects from the disruption are imposed, and vice-versa for the final 20 hours of the simulation where the disruption has ceased. We elect not to vary hazard duration for simplification purposes; the main goal of the simulation is not to achieve the highest possible fidelity but to be sufficient enough

to demonstrate the scalability of the machine learning procedures discussed later. As such, disaster scaling is more easily achieved by simply varying values for severity while keeping the disruption intervals constant. However, modeling the effects of hazard duration in conjunction with hard capacity limits is an alternative approach that we plan to explore in the future.

Overall, 24,000 disruption scenarios are simulated. The centroid locations of the hazards (longitude and latitude coordinates within the study area) are sampled via stratified randomization to ensure spatial coverage, and the radius of the hazard is randomly sampled from 1 to 40 miles to capture a wide range of spatial dependencies between network components within the study area. As mentioned, severity is randomly sampled from the values 2, 4, 8, 16, and 32 to help incorporate effects of "minor" to "severe" disruptions. Given these input variables, the primary outputs of the simulations are the total changes in travel times and travel distances aggregated across all routes relative to baseline conditions.

It should be noted that we impose the same schedule of dispatched OD routes for when we run the simulation with a disruption as we do for the baseline scenario (i.e., no hazard present). It is likely that during a large-scale disruption network demand would decreased due to reduced commercial activity and expected transportation issues. However, in the interest of maintaining consistent conditions to allow for more direct comparison of impacts between hazard scenarios and that of the baseline case, we assume there are no changes in travel demand. This assumption can of course be altered or explored further in future studies.

2.3 Gradient Boosting Machines

We then train gradient boosting machines (GBMs) on the 24,000 simulation runs, treating each scenario as an observed sample. Specifically, the GBMs take as inputs a hazard's size (radius), severity (edge traversal penalty), and centroid location (latitude and longitude coordinates) to predict changes in total travel times and distances caused by the given disruption. The GBMs act as surrogate models that allow us to interpolate outcomes with low computational effort compared to the simulation. Unlike typical, parametric regressions, such as what Haghighi et al. (2018) used for their surrogate model, GBMs do not require simplifying assumptions of variable interactions; instead, the GBMs non-parametrically learn these interactions (Chipman et al., 2010; Friedman, 2001). This feature is beneficial when trying

to model the complex, and often unknown, spatial dependencies between network components that manifest during area-spanning disruptions.

GBMs are a subset of a class of powerful machine learning algorithms called ensemble decision trees, which include bagging (Breiman, 1996), random forests (Breiman, 1996), and boosting (Freund and Schapire, 1999; Friedman, 2001; T. Chen and Guestrin, 2016). Each of these techniques uses a different non-parametric approach to fit linear combinations of decision trees to predict regression or classification responses, depending if the outcome of interest is continuous or categorical respectively. Boosting algorithms in particular have achieved prominent success in a variety of research applications and machine learning contests (Chipman et al., 2010; Natekin and Knoll, 2013).

Boosting algorithms use a sequence of small decision trees to progressively fit variability in the data that is not accounted for by earlier trees in the sequence (Chipman et al., 2010; Friedman, 2001). These algorithms take slow, incremental steps toward modeling the data, which helps the model avoid over-fitting the data and improve predictive performance (Friedman, 2001). This gradual procedure also enables the boosting algorithms to naturally incorporate complex variable interactions and additive effects in its formulation, which is crucial for effectively capturing spatial correlations in our freight network that arise during area-spanning disruptions (Chipman et al., 2010).

Of the several boosting algorithms noted in the literature, we predominantly use GBM in our analysis (Friedman, 2001). With GBMs, a successive tree is added to the ensemble by minimizing the gradient of a loss function, specified by the modeler, with respect to the variability in the data not accounted for by the current ensemble (Natekin and Knoll, 2013). Each successive tree contributes the same weight to the ensemble's overall prediction, controlled by a learning rate parameter. Related, an enhanced version of GBM, XGBoost, features second order derivatives in its loss functions and uses regularization techniques to construct trees (T. Chen and Guestrin, 2016). Both GBM and XGBoost have proven to be powerful techniques in applied and academic research (Natekin and Knoll, 2013). We attempted both in our analysis; however, although we were able to achieve slightly more accurate predictions in test samples with XGBoost compared to GBM ($\sim 4.5\%$ improvement in RMSE), deriving PDPs via *R*'s **pdp** package took considerably more computational effort with the former ($>$ 24hrs run-time vs. $\sim$ 1hr). For this reason, we elected to use GBM instead of XGBoost.

When training our GBMs, the independent variables consist of a hazard's severity, radius, and centroid location. The response variables are the cumulative time delays (hours) and total increases in travel distance (kilometers) resulting from each disruption scenario. Since the response variables are continuous, our GBMs comprise regression trees with squared-error loss functions (Natekin and Knoll, 2013). Note, we log-transform the response variables prior to training the models to help improve predictions by making the squared-error loss function not as susceptible to being skewed by the exceedingly large and severe disruptions. We formulate one GBM for each response variable, using 80% of the 24,000 scenarios as training samples and the remaining 20% as test dataset. Using the training dataset, we tune the GBMs' hyper-parameters for number of trees, interaction depth, and learning rate via a parameter grid search that seeks to minimize 5-fold cross-validation out-of-sample errors. The final models consist of the best-fitting hyper-parameters as determined by these out-of-sample errors. Lastly, we assess the predictive performance of the final models on the test dataset. The Supplementary Materials include a complete dataset of the independent and response variables for all 24,000 scenarios as well as an accompanying R script that details the tuning procedures for the GBMs.

2.4 Partial Dependency Plots

Given that the GBMs sufficiently predict the test dataset, we then use these models to construct PDPs to provide a map that intuitively depicts spatial vulnerabilities in the network at a high degree of resolution. PDPs are functional depictions of the relationships between a set of specified input variables and model predictions while controlling for the effects of the other non-specified input variables (Friedman, 2001). They are analogous to parameter coefficients in multiple linear regression and can be viewed as a special case of what is more broadly known as average predictive comparisons for nonlinear and non-parametric models, such as neural networks, kernel-based methods, and tree ensembles (Gelman and Pardoe, 2007). In this application, they help researchers intuitively analyze the inner-workings of oftentimes abstruse non-parametric machine learning models.

The mathematical expression for the partial dependency function is given

in equation (2).

$$\hat{f}_{x_S}(x_S) = E_{x_C}\left[\hat{f}(x_S, x_C)\right] = \int \hat{f}(x_S, x_C) p(x_S) d(x_C) \tag{2}$$

Here, x_S represents the target set of features for which one wishes to know their effects on model predictions; x_C is the complement to this set and contains all other features whose values are marginalized over in the function. By marginalizing over the complimenting features, one obtains a function that depends only on the target features while accounting for their interactions with all the variables (Friedman, 2001). Here, our target variables consist of permutations of finely-spaced latitude and longitude coordinates (1/10th degree resolution), so the resulting PDPs depict a smooth geographical plot of the marginal effect a hazard's centroid location will have on total increased travel times and distances while controlling for hazard size, severity, and spatial dependencies embedded in the network.

The partial dependency function is estimated via Monte Carlo approximation, specified in equation (3).

$$\hat{f}_{x_S}(x_S) \approx \frac{1}{n} \sum_{i=1}^{n} \hat{f}(x_S, x_C^{(i)}) \tag{3}$$

Values for x_S are held constant as model predictions are evaluated at each row (i) of complimentary features $x_C^{(i)}$ in a dataset consisting of n total rows. Predictions are then averaged across the n sets of $x_C^{(i)}$ to obtain the partial dependency function for x_S. Values for x_S can simply be those observed in the data or specified by the modeler as a grid of points. It is assumed that x_S and $x_C^{(i)}$ are independent, for if not, average values may include data points that are unlikely, or impossible, to occur. Related, partial dependency functions are inherently interpolations of the data, so modelers should use caution when extrapolating results (Friedman, 2001).

For our analysis, x_S consists of 40,000 grid points of approximately 1/10 degree longitude and latitude coordinates that span the Middle Tennessee region (i.e., each resulting grid box corresponds to an area of roughly 1.5 km^2). In other words, each coordinate in the PDP depicts the marginal effect this particular location has on outcomes while controlling for the effects of the complimentary features present in all the simulated scenarios. We should also note that the latitude and longitude coordinates are independent from

the radius and severity variables, so we do not have to worry about the PDPs interpolating data points that are unlikely to occur.

Although PDPs can be derived for any statistical model ($\hat{f}_{x_S}$), we found great utility in pairing them with the GBMs because they are able to non-parametrically learn the spatial dependencies of the network during disruptions and accurate predict outcomes. These types of variable interactions would have to be explicitly defined, and likely over-simplified, in more traditional, parametric machine learning methods such as multiple linear regression. Although PDPs are commonly used to depict the marginal effects variables have on predictions, to the best of our knowledge, they have not been previously used to depict spatial vulnerabilities in networks. We use the **pdp** package in *R* to calculate the PDPs for our analysis, and the script for reproducing our plots can be found in the Supplementary Materials.

3 Results

3.1 Simulation Results

The primary outputs of the simulation are the aggregate changes in travel times and travel distances relative to baseline conditions. These aggregate metrics are tabulated from all the individual OD flows that occur during the simulation; each baseline OD flow has a corresponding counterpart in the disruption scenario. Figure 2 illustrates an example of one simulation (i.e., a disruption on the left and the baseline case on the right).

In this example, an active disturbance (pink shaded circle) occurs over downtown Nashville. At this point in the simulation run, most road traffic has diverted into more circuitous routes around the city and rail traffic has slowed down because of slower network segments inside the disruption. In the lower left time series plot, the pink section corresponds to the 20-hour time period when the disruption is active. Shortly after the disruption begins, the number of active routes (red dashed line) increases due to the added time required for routes to complete. The total distance traveled (blue line) increases slightly due to diverted routes.

We use another metric, called the *effective route distance*, as a diagnostic tool to help establish the validity of the simulations. This metric is calculated as the distance (miles) that would have been covered in a flow block's pure shortest path route, proportional to the amount of its disrupted route

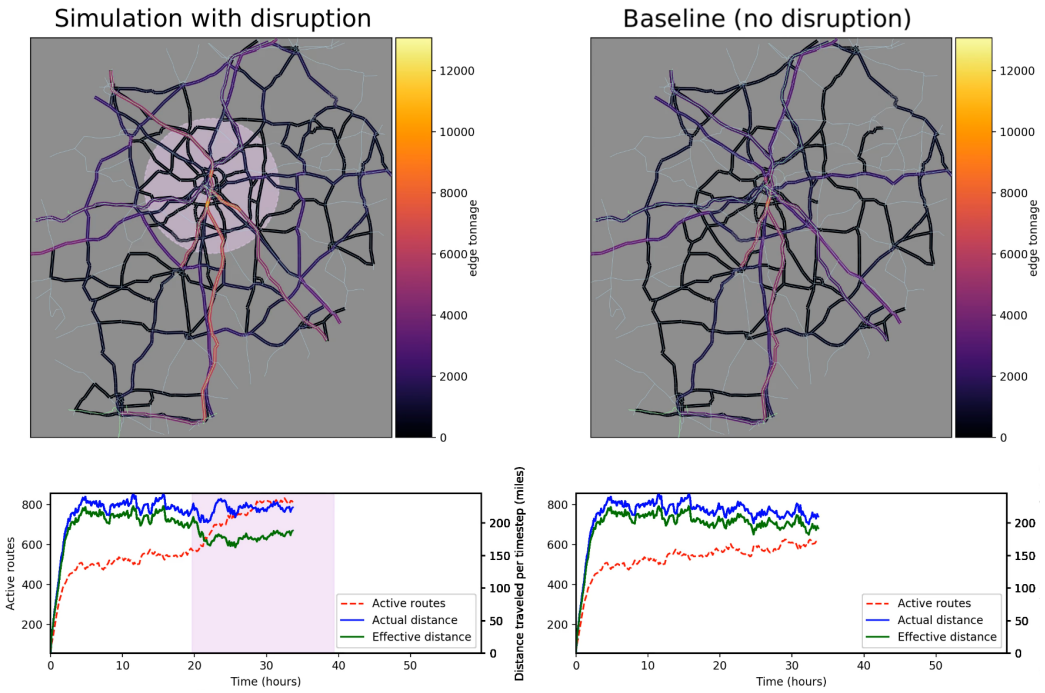


Figure 2: Map and time series plots taken mid-simulation at 33:40 simulation time for a disruption scenario (left) and the baseline case (right)

that it has covered. Conceptually, effective distance represents how much distance has been covered towards the destination; route detours still make progress towards the destination, but they cause effective distance covered to accumulate more slowly compared to a shortest path route. As seen with the green line in Figure 2, overall progress from the origin to destination transpires more slowly in the disruption case, even though higher-speed (less congested) detours might be chosen. Although this metric is useful for examining impacts of disruptions within simulation runs, it does not carry any meaning when we compare aggregated results between runs because effective route distance eventually comes back up to meet actual route distance. As such, we do not include it as one of the dependent variables for the GBMs.

Although it is not possible to directly validate simulated disruptions with real world data, results suggest that our simulations are reasonably intuitive. Routing behaviors respond immediately to disruptions and continue to change as traffic builds up in alternative routes, and the magnitude of these

route shifts and added congestion scale intuitively with severity. Additionally, impacts from disruptions align with expectations based on the topology of Middle Tennessee's road and rail network. For example, disruptions near downtown Nashville tend to result in traffic being diverted around the city via one of many alternative routes, while disruptions concentrated along links without nearby alternatives show a build-up of congestion along those routes (i.e., routes with no alternatives are forced to traverse the hazard).

3.2 Gradient Boosting Machine Results

The GBMs act as surrogate models that predict changes in aggregated travel times and distances given a simulated disruption's size, severity, and location. Figure 3 shows that the GBMs are able to accurately predict simulation outcomes. Note, they slightly overestimate effects of lower-impact disruptions and have higher variability when predicting them as well. However, predictions scale well as outcomes become more severe. The test sample root-mean-square errors (RMSEs) for the two models are 0.636 and 0.414 respectively. For comparison, the corresponding RMSEs for the null models (mean-only models) are 2.83 and 1.71.

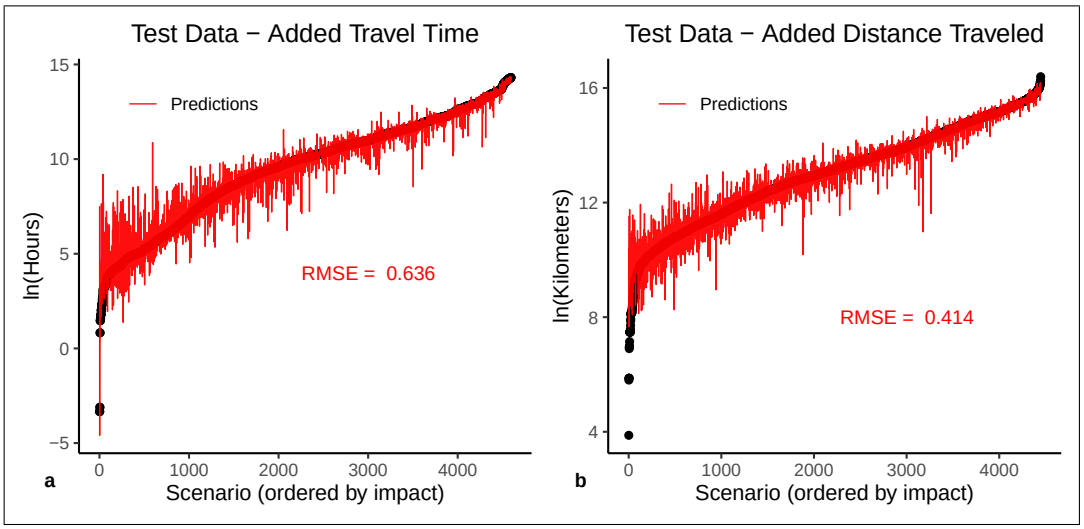


Figure 3: (a) GBM predictions vs. observations for total time delays; (b) GBM predictions vs. observations for total added travel distance

Table 1 displays the relative influence of the input variables. Relative

influence is the improvement, averaged across all decision trees, that each variable has on ensemble's predictions (Friedman, 2001). As Table 1 shows, the location-based variables (radius and geographic coordinates), are demonstrably the most influential factors in both GBMs. Additionally, found in the Supplementary Materials, both models' optimal hyper-parameters for interaction depth are large (30) relative to those of typical GBMs (∼3-5) (Wenxin, 2002). GBMs formed with larger variable interaction depths (i.e., larger trees) generally do not lead to notable, if any improvement in predictions compared to those formed with smaller interaction depths (i.e., smaller trees). However, given the intricate spatial dependencies we surmised were present in the network during area-spanning disruptions, it is not surprising that our GBMs preferred deep, complex interactions among the location-based variables. In other words, the GBMs slowly learn which combinations of network components spatially correlate with one another during disruptions to accurately predict impacts.

Table 1: Relative Influence of Variables

	Radius	Longitude	Latitude	Severity
GBM - Time	49.2%	24.8%	15.9%	10.0%
GBM - Dist	61.9%	17.8%	9.6%	10.7%

The PDPs derived from each GBM are presented in Figure 4. These plots depict the marginal effect a disruption's location (1/10th degree resolution) has on the outcome variable of interest while explicitly controlling for the effects of the complimentary features (i.e., radius and severity) and implicitly controlling for spatial correlations between coordinates that manifest during the area-spanning disruptions (i.e., the spatial dependencies the GBMs learned in order to accurately predict outcomes). As such, coordinates whose marginal effects predict the most severe outcomes can be interpreted as the most vulnerable areas with respect to that outcome.

For example, regarding time delays, the most significant impacts (i.e., most vulnerable areas) occur near major rail lines (yellow, Figure 4a). Nashville's rail network has few redundancies compared to its road network. As such, disruptions to rail lines result in lengthy time delays because only rail cars with intermodal capabilities can seek alternative road routes; the rest are

forced to traverse the disruption. However, with respect to increased travel distance, the most significant impacts occur near Nashville's city center, an area through which all major roadways pass (Figure 4b). In this situation, freight is redirected around the roads outside the city, so travel distance increases. Note, points outside the study area (black dashed-line, Figure 4) are extrapolations (i.e., the GBMs have not been exposed to these locations); as such, their values should be interpreted with caution.

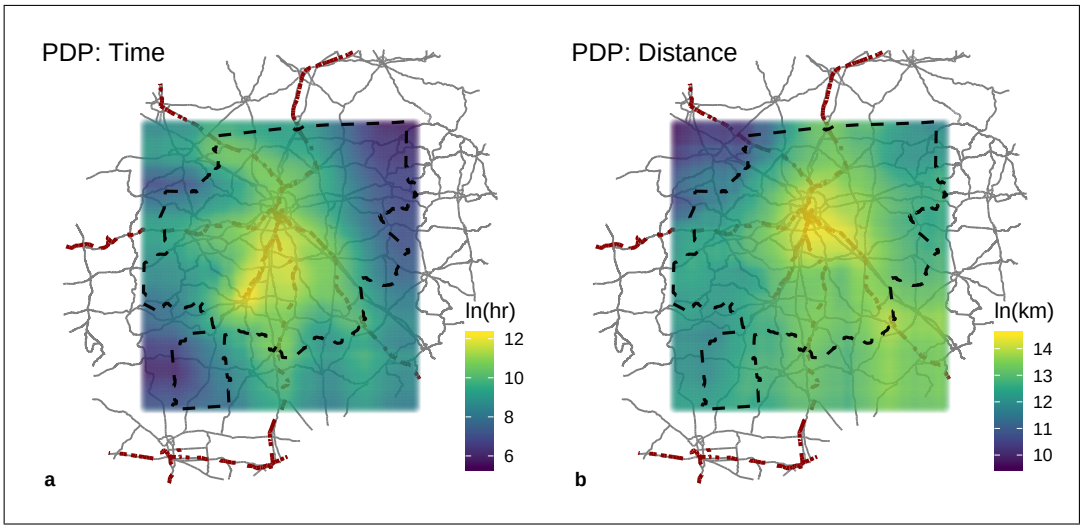


Figure 4: (a) PDP for total time delays - (black dashed-line) study area, (brown line) rail, (grey line) road; (b) PDP for total added travel distance

4 Discussion

The main benefit of using PDPs in this manner is that they provide a high-resolution depiction of the average impact a disruption's location has on outcomes while properly controlling for the effects of other input variables and spatial dependencies embedded in the network. As such, the PDPs function as an intuitive "heat-map" that reveals which areas in the network are vulnerable to area-spanning disruptions and/or may warrant opportunities for mitigation. For example, in Figure 4a, if decision makers' primary concern is to reduce time delays due to disruptions, focusing on rail lines in the southeast corner of Nashville's city center might be a prudent exercise, and

perhaps introducing more rail cars with intermodal capabilities would be an effective mitigation strategy.

Although PDPs can be paired with any statistical model, they worked particularly well with GBMs for this application. The GBMs were able to learn the complex spatial dependencies between components within the network and provide accurate predictions of simulation outcomes. In turn, the GBMs allowed us to interpolate all the scenarios necessary to create the PDPs at a fraction of the computational effort that would have been required to simulate them. Although achieving the same predictive power with parametric regression techniques is unlikely, other machine learning methods could serve as viable alternatives to the GBMs, such as other ensemble decision tree methods or neural networks. Additionally, as mentioned throughout the paper, the surrogate modeling and PDP framework can be applied to other traffic simulations and networks.

However, there are some drawbacks to our approach. One is that it can be difficult to isolate impacts to individual network components, should multiple components exist within given coordinate cell (here, 1/10th degree resolution or roughly 1.5 km^2). To remedy this issue, one could retrain the GBMs using the network components themselves as binary input variables, instead of longitude and latitude coordinates. Thus, the resulting PDPs would yield a diagram of the marginal effects from individual network segments, instead of coordinate points. Additionally, since we are not simulating every possible combination of multi-component disruptions, certain key scenarios may be left out of the training samples. The optimization approaches developed by Wang et al. (2016) and Xu et al. (2017) could be used to help establish bounds on potential impacts to help ensure the GBMs are exposed to the full range of impacts.

It should also be noted that our analysis does not currently consider the likelihood of hazards. We modeled the consequences of disruptions without considering their probabilities of occurrence because we were interested in revealing *vulnerabilities* to disruptions in transportation networks, not making overall risk assessments due to hazards (Haghighi et al., 2018; Sugishita and Yasuo Asakura, 2021). However, if we were to analyze the actual risks of hazards disrupting this network, the probabilistic nature of an event's severity, location, and shape could all be derived from empirical data and then inputted into the GBMs to predict expected impacts.

Additionally, were we to use this analysis for actual decision-making purposes, we would ideally have OD-level data for both road and rail flows to

ground the baseline conditions for the simulations and would also want to conduct sensitivity analyses on the routing logic and/or introduce uncertainty in routing decisions (He and Liu, 2012; Gao et al., 2016; Mahmassani, 1990). We would also want to link the simulations to regional economic models, so response variables could include fiscal impacts (Cho et al., 2001; Darayi, Barker, and Santos, 2017; Tatano and Tsuchiya, 2008).

5 Conclusion

We have demonstrated how researchers can use PDPs to intuitively depict spatial vulnerabilities to simultaneous, multi-component disruptions in transportation networks. The PDPs are derived from surrogate models that are trained on the results of a transportation disruption simulation where a subset of all possible link removal scenarios are analyzed. In particular, we found GBMs to be useful surrogate models for this application because they are a powerful class of non-parametric machine learning procedures capable of learning complex spatial interdependencies in networks. Given a disruption's location and other simulation input variables, the GBMs are able to accurately predict outcomes of our simulated scenarios, which allows us to evaluate outcomes of other link removal scenarios at a much lower computational cost than would be required to simulate them. In turn, the PDPs derived from the GBMs are able to interpolate and depict the marginal effects that a disruption's location (i.e., latitude and longitude coordinates) has on outcomes while controlling for the effects of other simulation variables and spatial correlations within the network that manifest during the area-spanning disruptions. Although we demonstrate our approach for Middle Tennessee's intermodal road and rail freight network, our framework can easily be applied to other networks.

Acknowledgments: This work was supported by TDOT through its State Planning and Research Program, project number 16SPR2-F7-035. Special thanks are given to Raquel Wright and Gary Baker at USDOT for helping us understand how to use their Waybill routing tool and to Jimmy Dobbins at Vanderbilt for teaching us how to properly use the Waybill data.

6 Appendix I - Modeling Middle Tennessee's Intermodal Freight Network

For constructing the road network, the Freight Analysis Framework (FAF) contains a variety of metrics pertaining to freight movement across states and metropolitan areas (Ford, 2017). At the time of this analysis, FAF4 was the most recent version, which bases metrics on commodity flows from the year 2012. The FAF metrics include but are not limited to daily freight volume per road segment, daily traffic volume per segment, and segment capacity. The road segments are defined over many intervals, each with identifying characteristics such as road name, speed limit, segment length, and number of lanes; each segment also provides a traversal time, which will be used later in routing across the network.

The FAF also houses geospatial data that includes a collection of georeferenced highways and roads across the US. This geospatial data is defined over the same roadway intervals as the commodity flow data, so we map the flow metrics to their corresponding edges to create a georeferenced road network dataset. Additionally, the FAF contains geospatial data of officially recognized commodity flow survey zones, with Middle Tennessee's zone (472) serving as the study area for our analysis. Within ArcGIS, we use the zone 472 shapefile with an added 50km buffer to clip the aforementioned roadway network dataset. Figure 1 depicts the buffered road network (green lines) and study area (gold). The buffer allows for more realistic routing behavior to take place during the simulation. For example, a truck whose origin and destination lie within zone 472 will reasonably travel outside that zone if it is quicker to do so. This type of situation commonly arises during disruptions. The size of the buffer is ultimately arbitrary, but given the size of the Middle Tennessee zone, 50km we consider to be a reasonable balance between allowing for flexibility in route choices while still focusing on the area of interest.

For the rail network, the Tennessee Department of Transportation (TDOT) supplied us with private Waybill data, records of O-D rail freight shipments, that passed through Tennessee for the year 2014. The records contain metrics such as but not limited to number of carloads, tons shipped, commodity type identifiers, and intermodal capabilities of rail cars. Publicly available samples of the Waybill data are available through the US Surface Transportation Board, should readers wish to look at the data structure.

Additionally, the US Department of Transportation (USDOT) manages publicly available, geospatial data of rail tracks and stations across the country. Similar to the road network, we map Waybill commodity flows to their corresponding segments in the geospatial data. However, since each Waybill entry corresponds to an individual shipment, records first need to be aggregated by unique legs of the rail network to get segment-level flow estimates. The USDOT has recognized the necessity of this task and developed a routing tool to work specifically with the Waybill data (Wright and Baker, 2017). We use this tool to create a map of expected flows of rail freight whose origin or destination lie within the Middle Tennessee area of interest (FAF Zone 472). This data also provides the network segment traversal time across the rail network.

7 Appendix II - Supplementary Materials

The dataset of all 24,000 disruption scenarios and the accompanying *R* scripts that were used to train the GBMs and produce the PDPs are available on Open Science Framework:
`https://osf.io/vezns/?view_only=725c89c810764be6865545f085316eef`

References

Asakura, Y (1999). "Reliability Measures of an Origin and Destination Pair in a Deteriorated Network with Variable Flow". English. In: *Transportation networks: recent methodological advances : selected proceedings of the 4th EURO Transportation Meeting*. Ed. by Michael G. H Bell. OCLC: 1010612652. United Kingdom: Emerland. ISBN: 978-0-08-043052-2.

Baghalian, Atefeh, Shabnam Rezapour, and Reza Zanjirani Farahani (2013). "Robust supply chain network design with service level against disruptions and demand uncertainties: A real-life case". In: *European Journal of Operational Research* 227.1, pp. 199–215. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2012.12.017.

Breiman, Leo (1996). "Bagging predictors". en. In: *Machine Learning* 24.2, pp. 123–140. ISSN: 0885-6125, 1573-0565. DOI: 10.1007/BF00058655.

Calatayud, Agustina, John Mangan, and Roberto Palacin (2017). "Vulnerability of international freight flows to shipping network disruptions: A multiplex network perspective". en. In: *Transportation Research Part E: Logistics and Transportation Review* 108, pp. 195–208. ISSN: 13665545. DOI: 10.1016/j.tre.2017.10.015.

Chen, Anthony et al. (1999). "A capacity related reliability for transportation networks". en. In: *Journal of Advanced Transportation* 33.2, pp. 183–200. ISSN: 2042-3195. DOI: 10.1002/atr.5670330207.

Chen, Tianqi and Carlos Guestrin (2016). "XGBoost: A Scalable Tree Boosting System". en. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*. San Francisco, California, USA: ACM Press, pp. 785–794. ISBN: 978-1-4503-4232-2. DOI: 10.1145/2939672.2939785.

Chipman, Hugh A., Edward I. George, and Robert E. McCulloch (2010). "BART: Bayesian additive regression trees". EN. In: *The Annals of Applied Statistics* 4.1, pp. 266–298. ISSN: 1932-6157, 1941-7330. DOI: 10.1214/09-AOAS285.

Cho, Sungbin et al. (2001). "Integrating Transportation Network and Regional Economic Models to Estimate the Costs of a Large Urban Earthquake". en. In: *Journal of Regional Science* 41.1, pp. 39–65. ISSN: 1467-9787. DOI: 10.1111/0022-4146.00206.

Chow, Joseph Y. J. and Amelia C. Regan (2014). "A surrogate-based multiobjective metaheuristic and network degradation simulation model for

robust toll pricing". en. In: *Optimization and Engineering* 15.1, pp. 137–165. ISSN: 1389-4420, 1573-2924. DOI: 10.1007/s11081-013-9227-5.

Darayi, Mohamad, Kash Barker, and Charles D. Nicholson (2019). "A multi-industry economic impact perspective on adaptive capacity planning in a freight transportation network". en. In: *International Journal of Production Economics* 208, pp. 356–368. ISSN: 09255273. DOI: 10.1016/j.ijpe.2018.12.008.

Darayi, Mohamad, Kash Barker, and Joost R. Santos (2017). "Component Importance Measures for Multi-Industry Vulnerability of a Freight Transportation Network". en. In: *Networks and Spatial Economics*, pp. 1–26. ISSN: 1566-113X, 1572-9427. DOI: 10.1007/s11067-017-9359-9.

Dingler, Mark et al. (2010). "Determining the causes of train delay". In: *AREMA Annual Conference Proceedings*.

Ford, Chester (2017). *Freight Analysis Framework (FAF)*.

Fotakis, Dimitris et al. (2002). "The structure and complexity of Nash equilibria for a selfish routing game". In: *International Colloquium on Automata, Languages, and Programming*. Springer, pp. 123–134.

Freiria, Susana, Alexandre O. Tavares, and Rui Pedro Julião (2015). "The Multiscale Importance of Road Segments in a Network Disruption Scenario: A Risk-Based Approach". en. In: *Risk Analysis* 35.3, pp. 484–500. ISSN: 1539-6924. DOI: 10.1111/risa.12280.

Freund, Yoav and Robert E Schapire (1999). "A Short Introduction to Boosting". en. In: *Journal of Japanese Society for Artificial Intelligence* 14.5, pp. 771–780.

Friedman, Jerome H. (2001). "Greedy function approximation: A gradient boosting machine." en. In: *The Annals of Statistics* 29.5, pp. 1189–1232. ISSN: 0090-5364, 2168-8966. DOI: 10.1214/aos/1013203451.

Gao, Bo, Ronghui Zhang, and Xiaoming Lou (2016). "Modeling Day-to-day Flow Dynamics on Degradable Transport Network". In: *PLoS ONE* 11.12. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0168241.

Gedik, Ridvan et al. (2014). "Vulnerability assessment and re-routing of freight trains under disruptions: A coal supply chain network application". en. In: *Transportation Research Part E: Logistics and Transportation Review* 71, pp. 45–57. ISSN: 13665545. DOI: 10.1016/j.tre.2014.06.017.

Gelman, Andrew and Iain Pardoe (2007). "Average Predictive Comparisons for Models with Nonlinearity, Interactions, and Variance Components". en. In: *Sociological Methodology* 37.1, pp. 23–51. ISSN: 0081-1750, 1467-9531. DOI: 10.1111/j.1467-9531.2007.00181.x.

Haghighi, Nima et al. (2018). "A Multi-Scenario Probabilistic Simulation Approach for Critical Transportation Network Risk Assessment". en. In: *Networks and Spatial Economics* 18.1, pp. 181–203. ISSN: 1566-113X, 1572-9427. DOI: 10.1007/s11067-018-9392-3.

Hartig, F. et al. (2014). "Technical Note: Approximate Bayesian parameterization of a process-based tropical forest model". en. In: *Biogeosciences* 11.4, pp. 1261–1272. ISSN: 1726-4189. DOI: 10.5194/bg-11-1261-2014.

He, Xiaozheng and Henry X. Liu (2012). "Modeling the day-to-day traffic evolution process after an unexpected network disruption". en. In: *Transportation Research Part B: Methodological* 46.1, pp. 50–71. ISSN: 01912615. DOI: 10.1016/j.trb.2011.07.012.

Ivanov, Dmitry (2020). "Predicting the impacts of epidemic outbreaks on global supply chains: A simulation-based analysis on the coronavirus outbreak (COVID-19/SARS-CoV-2) case". en. In: *Transportation Research Part E: Logistics and Transportation Review* 136, p. 101922. ISSN: 13665545. DOI: 10.1016/j.tre.2020.101922.

Jenelius, Erik and Lars-Göran Mattsson (2012). "Road network vulnerability analysis of area-covering disruptions: A grid-based approach with case study". en. In: *Transportation Research Part A: Policy and Practice* 46.5, pp. 746–760. ISSN: 09658564. DOI: 10.1016/j.tra.2012.02.003.

Kim, Yusoon, Yi-Su Chen, and Kevin Linderman (2015). "Supply network disruption and resilience: A network structural perspective". In: *Journal of Operations Management* 33-34.Supplement C, pp. 43–59. ISSN: 0272-6963. DOI: 10.1016/j.jom.2014.10.006.

Mahmassani, Hani (1990). "Dynamic models of commuter behavior: Experimental investigation and application to the analysis of planned traffic disruptions". In: *Transportation Research Part A: General* 24.6, pp. 465–484. ISSN: 0191-2607. DOI: 10.1016/0191-2607(90)90036-6.

Murray, Alan T. and Tony H. Grubesic (2007). "Overview of Reliability and Vulnerability in Critical Infrastructure". In: *Critical Infrastructure: Reliability and Vulnerability*. Ed. by Alan T. Murray and Tony H. Grubesic. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 1–8. ISBN: 978-3-540-68056-7. DOI: 10.1007/978-3-540-68056-7_1.

Murray-Tuite, Pamela and Hani Mahmassani (2004). "Methodology for Determining Vulnerable Links in a Transportation Network". en. In: *Transportation Research Record: Journal of the Transportation Research Board* 1882, pp. 88–96. ISSN: 0361-1981. DOI: 10.3141/1882-11.

Natekin, Alexey and Alois Knoll (2013). "Gradient boosting machines, a tutorial". en. In: *Frontiers in Neurorobotics* 7. ISSN: 1662-5218. DOI: 10.3389/fnbot.2013.00021.

Organization, Nashville Area Metropolitan Planning (2010). *Nashville Regional Freight & Goods Movement Study, Phase II: Existing Trends and Conditions*.

Patterson, S.A. and G.E. Apostolakis (2007). "Identification of critical locations across multiple infrastructures for terrorist actions". en. In: *Reliability Engineering & System Safety* 92.9, pp. 1183–1203. ISSN: 09518320. DOI: 10.1016/j.ress.2006.08.004.

Peng, Min, Yi Peng, and Hong Chen (2014). "Post-seismic supply chain risk management: A system dynamics disruption analysis approach for inventory and logistics planning". In: *Computers & Operations Research*. Multiple Criteria Decision Making in Emergency Management 42.Supplement C, pp. 14–24. ISSN: 0305-0548. DOI: 10.1016/j.cor.2013.03.003.

Saw, Krishna, B K Katti, and G Joshi (2014). "Literature Review of Traffic Assignment: Static and Dynamic". en. In: p. 9.

Snyder, Lawrence V. and Mark S. Daskin (2007). "Models for Reliable Supply Chain Network Design". en. In: *Critical Infrastructure*. Advances in Spatial Science. Springer, Berlin, Heidelberg, pp. 257–289. ISBN: 978-3-540-68055-0. DOI: 10.1007/978-3-540-68056-7_13.

Sohn, Jungyul (2006). "Evaluating the significance of highway network links under the flood damage: An accessibility approach". In: *Transportation Research Part A: Policy and Practice* 40.6, pp. 491–506. ISSN: 0965-8564. DOI: 10.1016/j.tra.2005.08.006.

Sugishita, Kashin and Yasuo Asakura (2021). "Vulnerability studies in the fields of transportation and complex networks: a citation network analysis". en. In: *Public Transport* 13.1, pp. 1–34. ISSN: 1866-749X, 1613-7159. DOI: 10.1007/s12469-020-00247-9.

Sullivan, James, Lisa Aultman-Hall, and David Novak (2009). "A review of current practice in network disruption analysis and an assessment of the ability to account for isolating links in transportation networks". en. In: *Transportation Letters* 1.4, pp. 271–280. ISSN: 1942-7867, 1942-7875. DOI: 10.3328/TL.2009.01.04.271-280.

Systematics, Cambridge (2007). "National rail freight infrastructure capacity and investment study". In: *Cambridge Systematics, Cambridge*.

Tatano, Hirokazu and Satoshi Tsuchiya (2008). "A framework for economic loss estimation due to seismic transportation network disruption: a spatial

computable general equilibrium approach". en. In: *Natural Hazards* 44.2, pp. 253–265. ISSN: 0921-030X, 1573-0840. DOI: 10.1007/s11069-007-9151-0.

UNDP (2011). *TOWARDS HUMAN RESILIENCE: Sustaining Mdg Progress in an Age of Economic Uncertainty*. Tech. rep.

USDOT (2019). *Freight Shipments by Mode — Bureau of Transportation Statistics*.

Wang, David Z. W. et al. (2016). "Identification of critical combination of vulnerable links in transportation networks – a global optimisation approach". In: *Transportmetrica A: Transport Science* 12.4. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/23249935.2015.1137373, pp. 346–365. ISSN: 2324-9935. DOI: 10.1080/23249935.2015.1137373.

Wenxin, Jiang (2002). "On Weak Base Hypotheses and Their Implications for Boosting Regression and Classification". In: *The Annals of Statistics* 30.1. Publisher: Institute of Mathematical Statistics, pp. 51–73. ISSN: 00905364.

Whitman, Mackenzie et al. (2017). "Component Importance for Multi-Commodity Networks: Application in the Swedish Railway". eng. In: *Computers & Industrial Engineering* 112.doi:10.1016/j.cie.2017.08.004, pp. 274–288. ISSN: 0360-8352. DOI: http://dx.doi.org/10.1016/j.cie.2017.08.004.

Wright, Raquel and Gary Baker (2017). *Modernizing FRA's Rail Freight Routing and Analysis Tools*.

Xu, Xiangdong, Anthony Chen, and Chao Yang (2017). "An Optimization Approach for Deriving Upper and Lower Bounds of Transportation Network Vulnerability under Simultaneous Disruptions of Multiple Links". en. In: *Transportation Research Procedia* 23, pp. 645–663. ISSN: 23521465. DOI: 10.1016/j.trpro.2017.05.036.