Analytically Comparing Disaster Resilience across Multiple Dimensions

Christopher W. Zobel *
Department of Business Information Technology
Pamplin College of Business
Virginia Tech, Blacksburg, VA 24061-0235
email: czobel@vt.edu

Milad Baghersad
Department of Operations and Supply Chain Management
Monte Ahuja College of Business
Cleveland State University, Cleveland, OH 44115
email: m.baghersad@csuohio.edu

ABSTRACT

It is important to compare the resilience of complex human systems to different types of disasters, in order to assess their inherent vulnerabilities and take appropriate actions to strengthen them. Resilient behavior can be complicated and multi-dimensional, however, and one must be able to characterize the different ways in which that resilience actually exhibits itself in practice. With this in mind, this paper discusses creating a multi-dimensional indicator for the resilience of a complex human system, and it explores an approach for visualizing and analyzing the relationships between each of the individual resilience dimensions. Because decision makers may differ on the relative contribution of the different dimensions to overall resilience, the paper further discusses the issue of weighting the different dimension values and the impacts that such a weighting scheme can have on the relative ranking of different scenarios. We illustrate the ability to characterize the complexity of multi-dimensional resilience by analyzing an empirical data set that measures the relative resilience of the New York metropolitan area to seven different natural disasters between 2010 and 2012.

Keywords: Decision Support Systems; Disaster Resilience; 311 Services; Indicators

* Corresponding author

1. Introduction

Natural disasters are increasing in frequency and intensity [1–3], and in the United States alone, disaster events are responsible for a growing cost of $57 billion per year on average [4]. Consequently, the ability to withstand and recover from the impacts of such events is being looked at in the context of a variety of different critical human systems, including critical infrastructure [5,6], supply chains and distribution systems [7–10], and communities [11,12]. In particular, the United Nations Office for Disaster Risk Reduction [13] and the US government [14] have recently emphasized the importance of incorporating resilience into communities in order to combat the vulnerabilities of populations exposed to different events.

There are many different definitions of disaster resilience in the academic literature, and most of these definitions typically incorporate both the notion of resisting the negative impacts of an event and the notion of recovering from those impacts in a timely manner [15–18]. When discussing the disaster resilience of a given system, it is important to identify not only how resilient that system is overall but also what the system is resilient to and the way in which it exhibits that resilience. For example, the physical infrastructure of a coastal community in Southern California may be resilient to the direct impact of earthquakes because of strict local building codes, but those regulations may not also protect against landslides and thus blocked roadways could still compromise residents' access to food, water, and medical care for an extended period of time.

In their seminal work on quantitatively characterizing the disaster resilience of a community, as an example of a complex system, Bruneau et al. [19] address this multi-dimensionality in the context of four specific dimensions of resilience: *technical*, *organizational*, *social*, and *economic*. *Technical resilience* refers to the ability of component physical systems to maintain functionality in the face of a disaster; *organizational resilience* refers to the ability of responsible organizations to take actions that contribute to resilient outcomes in these circumstances; *social resilience* refers to the capacity for impacted groups and individuals to resist and recover from the impacts of a disaster event; and *economic resilience* refers to the ability to reduce both direct and indirect losses due to such an event [19]. Subsequent research contributions, like Cutter et al. [20]'s disaster resilience of place (DROP) model, expand on such categorizations of multi-dimensional behavior by identifying specific sets of quantitative indicator variables that could be used to analytically capture the different characteristics of resilience across its many dimensions. Cutter's work builds on the four dimensions suggested by Bruneau et al. [19] and suggests also adding dimensions such as *ecological resilience* (the resilience of the natural environment) and *community competence* (how well a community can maintain its population wellness and quality of life) [20].

In order to actually improve the resilience of a complex human system, we must have the ability to analytically compare the multiple dimensions of resilient behavior against some set of baseline values. The typical approach taken in the literature is to aggregate different indicators into a single resilience value by simply using an equally weighted linear combination of the values [11]. This is justified as a way to keep the process simple and transparent and to avoid concerns that weighting schemes based on indicator importance are not able to appropriately reflect decision makers' priorities [11]. Such aggregation is far from ideal, however, since it loses almost all of the richness of using multiple dimensions to characterize the system's resilience behavior.

To address this issue and maintain a more multi-dimensional perspective, this paper introduces a much more effective approach for quantitatively comparing complex resilience behaviors across multiple dimensions simultaneously. The new approach provides a major extension to one of the research directions that has grown from the work of Bruneau et al. [19], namely the formal analytical characterization of the tradeoffs between loss and recovery time that is associated with the concept of predicted resilience [21,22]. By extending this work into a multi-dimensional context, the significant richness of the decision environment can be better leveraged to understand the complexity of a disaster and thus to support much more effective mitigation and recovery efforts.

To illustrate the application of the approach, a new type of resilience related data is discussed: non-emergency 311 requests for municipal services. These structured requests cover a wide variety of different aspects of the relationship between a population and its municipality, and as such they can be used to indicate the different ways in which a disaster event impacts a community. This, in turn, allows the municipality to target future service operations so that they can more efficiently and effectively help their constituents recover from the effects of a large-scale hazard event, such as a major winter snowstorm. The discussion builds on a preliminary example presented in Zobel et al. [23], and specifically implements the approach in the context of five different indicators drawn from 311 data generated in New York City, as measured with respect to a series of natural hazards that impacted it between 2010 and 2012.

We begin by presenting an overview of the relevant disaster resilience literature. After motivating and discussing the approach for representing the multi-dimensionality, we look at the issue of weighting the relative contribution of the different resilience dimensions, in terms of support for comparing the multi-dimensional resilience characteristics of different systems. We then introduce the data and the specific case study underlying the discussion, and explain the application of the new approach. The paper concludes with an analysis and discussion of related research opportunities.

2. Background

Although significant progress has been made in improving disaster-related knowledge and technology over the past few decades, losses due to natural and man-made disasters are nevertheless continuing to increase [4]. Because it is extremely difficult, if not impossible, to avoid all of the negative impacts of these disaster events, a paradigm shift in governments' policies and in academic research has occurred over the past few years: a move from a focus on prevention to one on resilience, i.e. the ability to "bounce back" from the impacts of such disasters [24–26]. For example, in 2009, the U.S. Department of Homeland Security (DHS) started a new program, called the Regional Resiliency Assessment Program, which is led by the Office of Infrastructure Protection [14]. This program aims to evaluate the resiliency of critical infrastructure to disasters and to address a range of infrastructure resilience issues.

Over the past few years, different organizations and scholars have conceptualized disaster resilience in a number of different ways. For example, the U.S. National Academy of Sciences (NAS) defines disaster resilience as "the ability to prepare and plan for, absorb, recover from, and more successfully adapt to adverse events" [15]. Similarly, the United Nations Office for Disaster Risk Reduction (UNISDR) defines disaster resilience as "*the ability of a system, community or society exposed to hazards to resist, absorb, accommodate, adapt to, transform and recover from the effects of a hazard in a timely and efficient manner, including through the preservation and restoration of its essential basic structures and functions through risk management*" [16].

Most definitions of disaster resilience in the literature, such as that of the UNISDR, are general enough to be applied at different levels within a community, from the level of an individual household to that of an entire metropolitan area. This generalizability has allowed researchers to apply a number of different approaches and create a variety of different indices for quantifying such disaster resilience. At each level of analysis, these approaches can be divided into two main categories: capacity-oriented methods, and output-oriented methods. We explore each of these in the following discussion.

2.1. Capacity-oriented methods

Capacity-oriented methods evaluate the resiliency of a system in terms of its anticipated capacity for resisting a disruption and for enabling effective recovery. In capacity-oriented methods, researchers effectively assess the extent to which a system ***would*** react to a disaster event before any such events actually occur. These assessments therefore are not generally associated with a specific observed instance of any particular event. Two common types of capacity-oriented methods are survey-based methods and index-based methods.

Survey-based methods collect individuals' subjective evaluation of different resilience capacities. For example, Nguyen & James (2013) estimated households' resilience to floods using individual responses to a set of 10 questions. The questions were designed to assess three specific properties of households: (1)

households' ability to secure health, income, and food, and to evacuate during and after a flood; (2) households' confidence about the ability of their building to resist a severe flood; and (3) households' interest in learning and practicing the new flood-based practices. [28–33] are a few examples of other studies that applied a survey-based capacity-oriented approach for assessing the general level of resiliency of systems in different settings.

Index-based methods select a number of observable indicators, such as socio-economic variables, in order to characterize the anticipated capacity for a system to be resilient. The Resilience Capacity Index developed by Foster [34], for example, is a single resilience metric based on 12 equally weighted indicators (four indicators in three dimensions): (1) regional economic capacity (income equality, economic diversification, regional affordability, and business environment); (2) socio-demographic capacity (educational attainment, without disability, out of poverty, health-insured); and (3) community connectivity capacity (civic infrastructure, metropolitan stability, homeownership, and voter participation). Another well-known composite index is the baseline resilience index (BRIC) of Cutter et al. [11], which is based on 36 equally weighted indicators divided into five categories: social resilience, economic resilience, institutional resilience, infrastructure resilience, and community capital. The authors used publicly available data sources such as the United States Census and FEMA databases to calculate the values of the 36 indicators for different communities.

2.2. Output-oriented methods

Output-oriented methods measure the resiliency of a system in terms of its operational performance (functionality) after a specific (actual or simulated) disaster event. In contrast to capacity-oriented methods, therefore, output-oriented methods evaluate how well a system actually ***did*** react to a particular disruptive event. Such approaches often use a response curve to characterize the changing level of functionality of a system over time, such as in the example given by Figure 1. In this example, a sudden disruption causes the functionality of the impacted system, F(t), to drop from 100% to 50% at time t_0 before then recovering at a constant rate back to full functionality at time $t_0 + T$. The quantity T^* is a user-defined upper bound on the length of the recovery time.

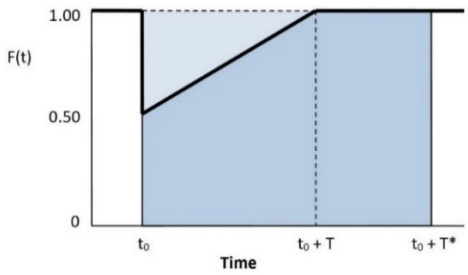

Fig. 1. Simple response curve (adapted from [35]**)**

Bruneau et al. [19] used the area above such a response curve (characterized as a *resilience triangle*) to define a measure for the loss of resilience in a system; a larger area thus corresponds to a less resilient system. Building on Bruneau et al.'s work, Zobel [21,36] subsequently developed a new measure of *predicted resilience*. This measure is based on the normalized area *under* the response curve, as a proportion of the total area under the curve that would have been realized if there had been no disruption. For example, for a response curve such as the one shown in Figure 1, the predicted resilience is given by the following formula:

$$R = 1 - \frac{XT}{2T^*} \tag{1}$$

where X is equal to the loss suffered at time t_0. Zobel & Khansa [22] generalized this measure by substituting the average loss over time ($\bar{X}$) into the equation, so that the resilience measure would apply to more general response curves without instantaneous loss or a requirement of linear recovery behavior:

$$R = 1 - \frac{\bar{X}T}{T^*} \tag{2}$$

One of the advantages of each of these approaches for calculating resilience is that they can be used to generate a series of parallel hyperbolic curves by which the tradeoffs between the amount of loss and the length of recovery time can easily be compared across multiple systems or across multiple possible outcomes for a single system [21,22,36,37]. Each curve represents a single level of resilience and reflects the fact that there are many different combinations of $\bar{X}$ and T that can result in the same total area beneath the response curve. These resilience curves, as illustrated in Figure 2, are the foundation for the development of our multi-dimensional resilience measure discussed below.

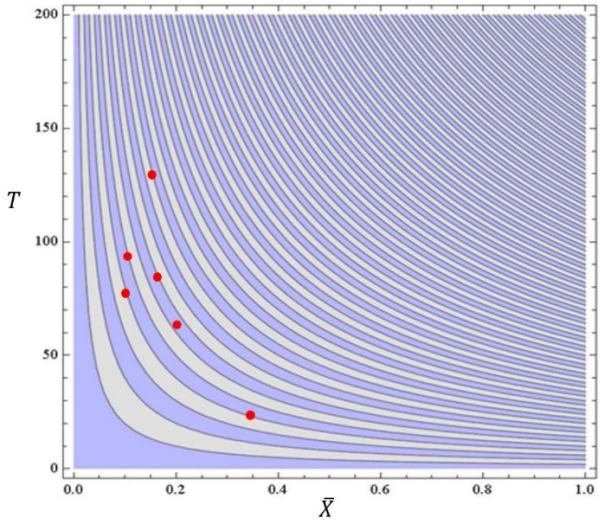

Fig. 2. Resilience curves with multiple resilience outcomes

The more general idea of using a response curve to quantify resilience has been widely extended and applied in different contexts, such as community resilience [38], health care facilities [39],

infrastructure systems and interdependent industry sectors [6,40–42], cyber security [22], transportation systems [43–45], and supply chains and organizations [46–48]. The concept of predicted resilience has also been extended to allow for adjusting the resilience curves based on decision makers' preferences [21], and to more accurately characterize different types of nonlinear recovery behavior [49].

In contrast with the multi-dimensional perspective of most capacity-oriented methods, output-oriented methods typically only evaluate the resiliency of a system based on a single representative dimension that is chosen to represent the functionality of that system. In most complex systems, however, such a single dimension cannot capture all the different aspects of the system's response that can be used to help characterize resilience. It is important, therefore, to be able to look at multiple dimensions at the same time, even in the case of an output-oriented approach, in order to better capture the system's overall ability to resist and then recover from a disaster. The following discussion seeks to address this need by generalizing the concept of predicted resilience, and its associated resilience curves, so that it can be applied to multiple dimensions simultaneously, and so that the relative contributions of each of those different dimensions to overall system resilience can be assessed both separately and in conjunction with the other dimensions being considered.

3. Methodology

The way that most capacity-oriented methods combine multiple indicators (and thus the different dimensions of resilience) into a single resilience value is by giving equal weight to each dimension [11]. Such equal weighting effectively implies, however, that all dimensions provide an equivalent amount of information about the system's complex behavior in response to a disaster event. In the context of a community's resilience, for instance, this implies that a broken streetlight after a disaster has the same level of importance as a downed tree in a public park. In reality, however, the relative importance of two such measures may very well depend on context. Individuals who live in a safer neighborhood and frequently visit the city's parks might assign more importance to the removal of the downed tree than they would to fixing a single street light. Individuals who live in higher crime areas, however, may be much more concerned about the impact of the broken street light on their family's safety. There thus can be an advantage to applying differing weights to different dimensions, in order to capture such differences in relative value.

As with the indicator-based capacity-oriented approaches, the multi-dimensional resilience metric developed below also provides the opportunity to combine individual dimensional indicators into a single overall resilience value, whether those indicators are equally or differentially weighted. In contrast to the indicators used in a capacity-oriented approach, however, if each indicator in the multi-dimensional metric is based on a metric such as predicted resilience, then it also provides the ability to explicitly visualize and

analytically characterize the tradeoffs between loss and recovery time within that particular dimension of the system's functionality. This also allows any weighting scheme to be applied to assessing overall loss and recovery across the different dimensions, and thus it supports providing a significantly richer characterization of the system's complex, overall resilience behavior. We develop the approach as follows:

Suppose, in general, that we have a complex system that can be characterized by n different dimensions of interest. We may then represent the general form of a combined n-dimensional measure of resilience to some event j as:

$$R_j = \sum_{i=1}^{n} w_{ij} R_{ij} \quad 0 \leq w_{ij} \leq 1 \tag{3}$$

where w_{ij} is the weight applied to the calculated resilience value for dimension i, given event j, and where, based on equation (2):

$$R_{ij} = R_{ij}\left(\bar{X}_{ij}, T_{ij}, T^*\right) = 1 - \frac{\bar{X}_{ij} T_{ij}}{T^*} \tag{4}$$

Here $\bar{X}_{ij}$ represents the average deviation from baseline functionality per unit of time, between the onset of the disaster event j and the end of an elapsed time period of T^* time units. T_{ij} then represents the number of time units, during that same time period, that were spent in a state of deviation due to event j.

Each calculated resilience value from equation (4) may be plotted on a set of resilience curves, as in Figure 2, through its $\bar{X}_{ij}$ and T_{ij} values. It turns out, however, that we cannot also plot the combined resilience value from equation (3) on these same curves by simply applying the same set of weights to the set of loss and recovery time values drawn from each individual dimension; this is due to the nonlinear nature of the calculation in equation (2). For example, as illustrated in Figure 3, if two different observations have the same level of resilience and we create a convex combination of the two values, then the resulting new observation should also have this same level of resilience, i.e., it should lie somewhere on the same resilience curve as the other points (the red line segment in the example of Figure 3). If we actually apply the same convex combination of weights independently to the corresponding $\bar{X}$ and T values, however, then the weighted $\bar{X}$ and T coordinates that are generated are those of an observation with a distinctly lower level of resilience. This issue must be addressed in order for us to correctly combine not

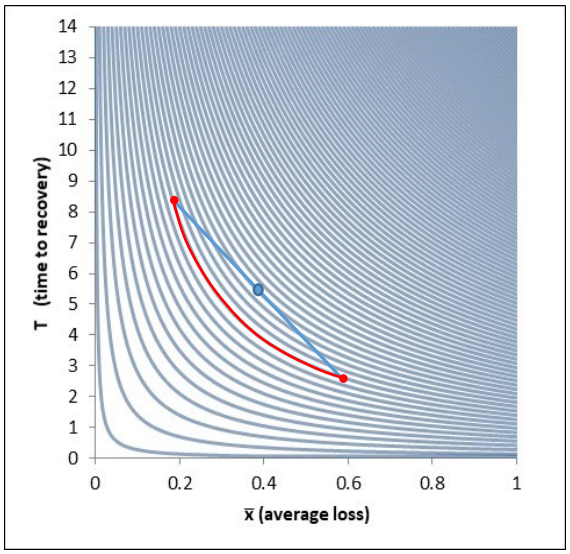

Fig. 3. Weighted average resilience vs. average $\bar{X}$ and T

just the resilience values but also the loss and recovery time values, and thus leverage the ability to characterize behavior using the resilience curves.

Our approach for accomplishing this recognizes that each individual observation of a resilience value is uniquely determined not just by its $\bar{X}$ and T values but also by the combination of R and the rescaled slope of the line that connects it with the origin: $T/\bar{X}$, from which those $\bar{X}$ and T values can be derived. Since the slope is non-linear, however, direct convex combinations of different slopes will be skewed towards larger values. We therefore substitute the angle of elevation from the $\bar{X}$ axis, α in place of the slope, in order to support creating convex combinations that have the desired linear behavior. This adjustment gives us the following result:

$$\alpha_j = \sum_{i=1}^{n} w_{ij}\alpha_{ij} \quad 0 \le w_{ij} \le 1 \tag{5}$$

where

$$\alpha_{ij} = \arctan\left(\frac{T_{ij}}{\bar{X}_{ij}T^*}\right) \tag{6}$$

and α_{ij} is based on a constant rescaling of the slope by 1/T*. The unique corresponding loss and recovery time values for combined resilience R_j with angle of elevation α_j are then given as follows:

$$T_j = T^*\left((1 - R_j)\tan(\alpha_j)\right)^{0.5} \tag{7}$$

$$\bar{X}_j = \frac{T^*(1-R_j)}{T_j} \tag{8}$$

since $(T_j/T^*) = (1/\bar{X}_j)*(1\text{-}R_j)$ implies that $(T_j/T^*)^2 = ((T_j/T^*)/\bar{X}_j)*(1\text{-}R_j)$.

This allows us to take the convex combination of any set of resilience observations, $\{(\bar{X}_{ij}, T_{ij})\}_{i=1\ldots n}$, and to generate a single weighted resilience value, R_j, together with a corresponding resilience observation,

$(\bar{X}_j, T_j)$, that can then be plotted in the context of the resilience curves. With regards to the example in Figure 3, because this convex combination is actually performed on the equivalent set of observations in R and α, $\{(R_{ij}, \alpha_{ij})\}_{i=1...n}$, the result of combining any two observations with the same level of resilience will always be a third observation with that same resilience value.

This new ability to create a differentially weighted combination of resilience observations allows us to support a variety of different types of analyses, such as comparing the overall resilience of a single system across different types of disaster events or comparing several different systems with respect to how resilient they were to the same disaster event. Given the $\bar{X}$ and T values for each constructed resilience observation, we can also extend the analyses to include a comparison of the average loss values and the average recovery time behavior. Furthermore, because we also have the $(\bar{X}, T)$ coordinates for each individual dimension of the system or systems, we can compare system performance on the basis of how each individual dimension reacts to a disaster event.

The following case study illustrates each of these types of analysis, and it provides a relevant context within which we can further extend our discussion to the use of non-equal weighting schemes in equations (3) and (5). The data set being discussed was initially described in Zobel et al. [50], but it was only analyzed with respect to the individual dimensions and a very simple convex combination of loss and recovery time. The discussion below therefore makes several important contributions to the literature, both in the development of the multi-dimensional approach and its support for differential weighting schemes, and in the application and analysis of that approach within the specific context of the case study.

4. Application: Non-emergency 311 requests for municipal services

One of the main responsibilities of a municipality is offering an appropriate level of public services to its citizens, such as maintaining roadways or providing access to the water and sewer systems. Any interruption in providing such public services can potentially disrupt lives and lead to significant economic losses. Natural disasters can affect the ability of municipal agencies to provide public services both because of the direct impacts from the disasters and because of an increasing number of public service requests from citizens. From a disaster resilience standpoint, changes in the interactions between the population and the municipality during and after a disaster event can indicate the relative impact of that disaster on different parts of the community, from both a geographic and a socio-economic perspective. With a better understanding of these impacts, a municipality can make more informed choices about how to target their limited resources as they prepare for and respond to future disasters.

Non-emergency 311 service requests provide an interesting example of a data set that can capture the interactions between a municipality and its citizens, in the context of such service provision [51]. Many large cities in the United States make their 311 requests publicly available: New York City, for example,

makes all of its 311 data from 2004 to the present available through its Open Data initiative (https://nycopendata.socrata.com/). Each individual service request in this NYC 311 data set is associated with a number of different attributes, a selection of which are presented in Table 1, and the data set includes more than 100 different complaint types that are directed to different agencies.

Our case study is based on this services supply chain data set, with a focus on service requests made between 2010 and 2012, and with consideration of five specific complaint types: Damaged Tree, Traffic Signal Condition, General Construction, Street Light Condition, and Blocked Driveway. These complaint types were selected, in particular, not only because they were all among the top types of complaint recorded in 2012 but also because they give a broad characterization of the community's health from a number of different perspectives (see Table 2). We therefore would expect their combined behavior, in response to a disaster event, to provide a reasonable approximation of the city's overall ability to resist and then recover from that event from a services perspective. Zobel et al. [50] provides more details about the process of retrieving and processing this data, as well as a more complete discussion of the set of agencies and the different types of complaints registered in the system.

Table 1: Selected 311 service request attributes

Attribute	Description	Attribute	Description
Created Date	Date and time the record was created	Incident Address	Street address of incident location
Closed Date	Date and time the record was closed	City	City of incident location
Agency Name	Specific agency name	Borough	Borough of incident location
Complaint Type	Category of complaint type	Due Date	Date and time the request is due
Descriptor	Detailed description of complaint	Latitude	Latitude of incident location
Incident Zip	Zip code of incident location	Longitude	Longitude of incident location

Table 2: Relations between selected complaint types and different aspects of a community

Category	Dimension	Complaint type
Critical Infrastructure	Transportation	Traffic signals
	Healthcare and public health	Streetlights
	Commercial facilities	General construction
Daily activities	Ability to get to work	Blocked driveway
Public safety	Quality of life / Property value	Damaged tree

Table 3: The seven disaster events selected between 2010 and 2012

Event ID	Event description	Date of Event
Event 1	2010 Nor'Easter	13-14 Mar. 2010
Event 2	Brooklyn / Queens tornadoes	16 Sep. 2010
Event 3	N. American Blizzard	25-27 Dec. 2010
Event 4	Hurricane Irene	28 Aug. 2011
Event 5	Major Snowstorm	31 Oct. 2011
Event 6	Hurricane Sandy	29 Oct. 2012
Event 7	2012 Nor'Easter	7 Nov. 2012

4.1. Preliminary analysis

A total of seven major disaster events that occurred in New York City between 2010 and 2012 were selected for analysis; some of these disaster events were localized, like the tornadoes that hit Brooklyn and the Bronx in 2010, and some of them, like Hurricanes Irene and Sandy, impacted the entire New York City metropolitan area. Table 3 presents these seven disaster events along with the date that each of them impacted the city. In order to analytically assess the impacts of these disasters, we calculated the number of daily service requests for each complaint type over the entire two-year time period and generated a set of corresponding time series response curves.

Figure 4 illustrates the particular curve that was constructed from the daily number of requests in the *Damaged Tree* dimension. With the exception of Event 3 (the 2010 North American Blizzard), each event appears to have resulted in a significant temporary increase in the daily number of *Damaged Tree* service requests. This type of behavior is reflected in the other dimensions as well, although in some cases the events instead appear to decrease the number of service requests. For example, Figure 5 indicates that

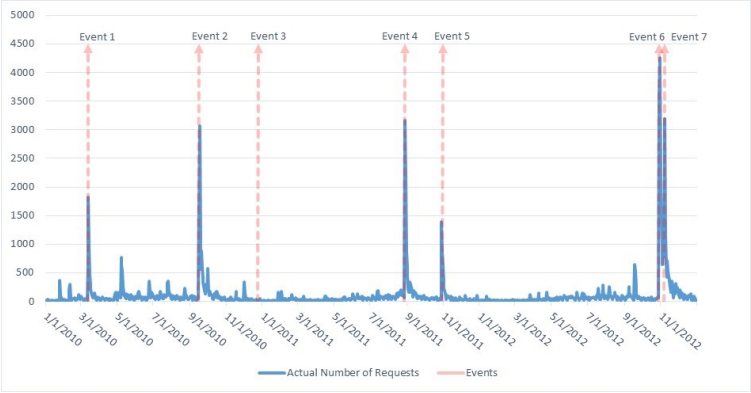

Fig. 4. Number of *Damaged Tree* service requests

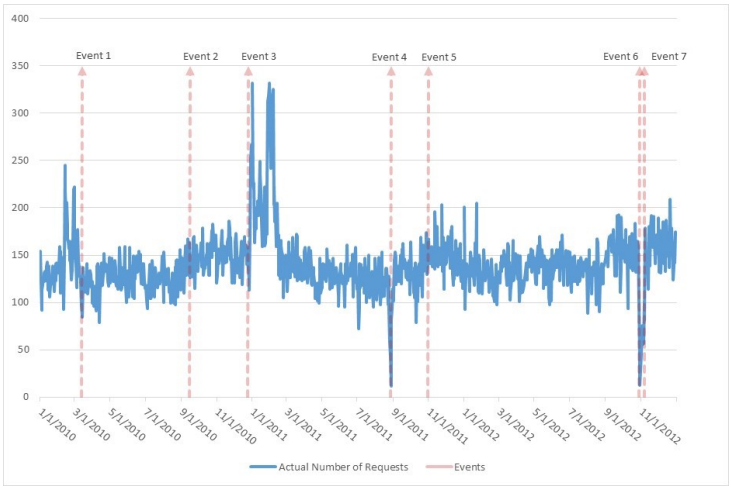

Fig. 5. Number of *Blocked Driveway* service requests

the number of *Blocked Driveway* service requests significantly decreased during both Hurricane Irene (Event 4) and Hurricane Sandy (Event 6), most likely because fewer vehicles were on the roads during these widespread events.

In order to more precisely measure the relative impact of each of these seven events, we used the three years' worth of data prior to 2010 to construct 95% prediction intervals for the normal daily number of service requests for each complaint type, using a quantile regression forest approach [52]. We then calculated the difference between the actual number of service requests and the appropriate corresponding upper or lower bound of the prediction interval to achieve an estimate of the net daily change in request volume due to the events. Using a two week time period after each disaster occurred (T^*=14 days) as a threshold, these net daily changes then were used to calculate the average net daily change associated with each dimension, ($\bar{X}_{ij}$), together with the total amount of time spent in a net state of change, (T_{ij}). This was done for each combination of complaint type $i \in \{1,\ldots,5\}$ and event $j \in \{1,\ldots,7\}$.

These paired loss and recovery time values were then plotted on a set of resilience curves to provide an indication of the different dimensional resilience values (see Figure 6a). Generally speaking, Figure 6a indicates that New York City exhibited less resilience to both Hurricane Sandy and the 2010 Blizzard in most, but not all, of the different resilience dimensions. The overall behavior becomes more clear, however, when we use equations (3) and (5), with nominally equal weights for each dimension, to generate a single combined value for the city's resilience to each particular event (see Figure 6b). Now it more easily can be seen that the city was most resilient to the 2011 Snowstorm, and that the tornadoes in 2010 resulted in the least overall amount of average loss. It is also easy to see that the loss due to the tornadoes was offset by a long enough recovery time that the city ended up exhibiting slightly less resilience to the tornadoes than it did to the Snowstorm. Furthermore, the resilience associated with the Snowstorm has the same relative

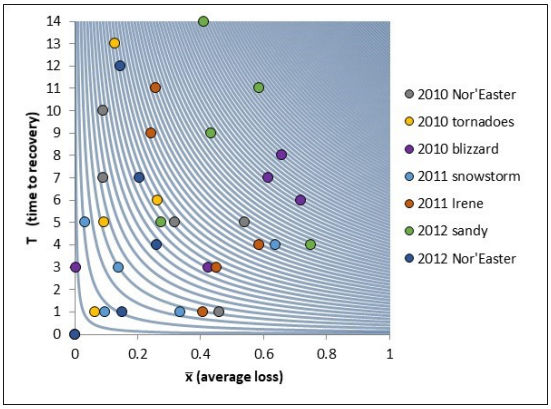

Fig. 6a. Individual dimensional resilience measures for each event

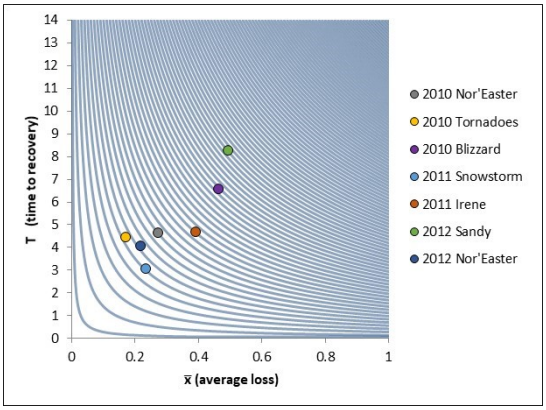

Fig. 6b. Average resilience across all dimensions for each event

ratio of average loss to recovery time as that for Hurricane Irene, indicating that even though Irene had a larger overall impact, the population responded to both events in a similar manner overall.

4.2. Differential weighting of multiple dimensions

The weighting approach developed in equations (3) and (5) will apply whether or not the different dimensional weights have equal values. If we allow for differing values, however, then we must extend the discussion a bit further. Assigning different weights to different dimensions implies that a decision maker is motivated to emphasize one dimension over another, as discussed above, but it does not necessarily follow that they are certain about the actual extent to which it should be emphasized. Because changing the weights on the different dimensions will change the overall resilience value for the system, however, a slight change in the amount of emphasis given to a particular dimension could potentially cause the relative ranking of the resilience of two different systems to be reversed, based on the new set of weights. Although this is not necessarily a problem, in and of itself, we must recognize and quantify such behavior in order to support more informed decision making. We therefore extend our discussion to consider the range of possible values that a differentially-weighted multi-dimensional system might take on under different circumstances.

Given the combined measures of resilience and angle of elevation from equations (3) and (5), along with a specific event, j, we may characterize the range of the impact that different weighting schemes can have on the combined resilience value by generating the convex hull for the original set of observations: $\{(R_{ij}, \alpha_{ij})\}_{i=1\ldots n}$. There are well-known approaches available for doing this analytically [53], and the end result can be represented as a two-dimensional boundary within which all convex combinations of those vertices, and thus all possible weighted resilience values, must lie.

In order to more easily analyze the relationship between the combined resilience values, and to assess the inherent tradeoffs between loss and recovery time, we translate this initial linear boundary onto

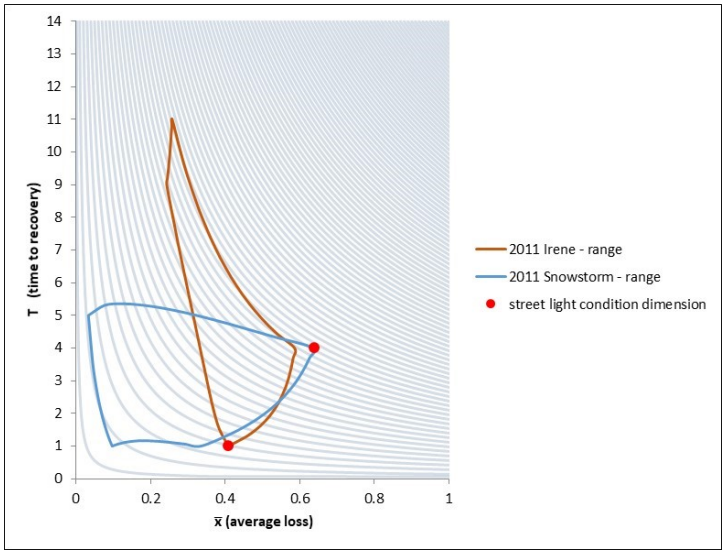

Fig. 7. Convex hulls of weighted resilience values for Irene and the 2011 Snowstorm

the nonlinear resilience curves in $\bar{X}$ and T by using the unique relationship between each (R_j, α_j) and its corresponding $(\bar{X}_j, T_j)$, for any set of weight values. This establishes an important opportunity to directly compare the response of a community to different events and to characterize the range of "possible" resilience values that could result from weighting the different dimensions. Figure 7 illustrates this for two of the events that struck New York City in 2011, Hurricane Irene and the major Snowstorm, and it clearly shows the nonlinear nature of the boundaries that results from the translation onto the $\bar{X}$-T curves.

Comparing convex hulls such as these can indicate, to a certain extent, which event has an overall resilience value that tends towards less loss and/or less recovery time across different possible weighting schemes. It is important to note, however, that the vertices (i.e., the individual resilience dimensions) on which the convex hulls are based will not necessarily be oriented in the same way. For example, giving more weight to the vertex representing resilience in the *street light condition* will increase the overall combined resilience in the case of Irene, but at the same time, it will decrease the overall resilience associated with the Snowstorm. The indicator variable for that particular dimension therefore contributes to the overall resilience of each event very differently.

One of the distinct advantages of using the resilience curves to represent the range of resilience behaviors under different weighting schemes is that we can examine multiple outcomes simultaneously. By generating a lattice of points based on varying the relative weights applied to the different dimensions, each distinct combination of weights will result in a set of related points, with each individual point associated with exactly one of the convex hulls being analyzed. By then comparing the points associated with a given set of weights, we can judge, for that set of weights and the corresponding priorities that they

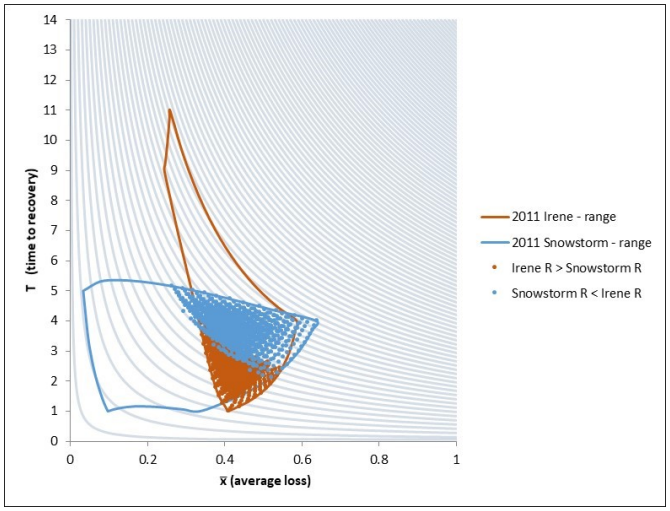

Fig. 8. Paired weighted observations representing more resilience to Irene than the Snowstorm

represent across the different dimensions, to which event the system was more resilient and to which event it had less overall loss or a quicker overall recovery time.

Because we are able to analytically compare these individual points across events, we can identify the sets of weights that satisfy certain conditions. For example, Figure 8 illustrates the set of points for which the city was more resilient to Hurricane Irene than to the 2011 Snowstorm when the weights are systematically varied in increments of 0.01. This reflects the fact that even though the two hulls have a relatively large overlap in the area that they cover, only 1052 out of 10610 total generated observations, or 11.8% of the total weight combinations, actually correspond to this situation. Similarly, if one compares the paired observations between the two events with respect to their average loss values, rather than resilience, only 6.8% of the points correspond to Irene having less average loss under the same weighting scheme, despite the extent of the overlap between the two hulls. In the same way, only 16.7% of the total combinations lead to a quicker recovery time in the case of Hurricane Irene.

From the standpoint of better understanding the implications of decision makers' preferences among different measures, we may also use these results to determine the relative magnitude of weights on the different dimensions that actually lead to the observed behaviors. For example, Figure 9 gives box plots ({min value, Q1, median, Q3, max value}) for the combined range of each set of sample weights that correspond to more resilience to Irene than to the Snowstorm. These results imply, in this particular case, that the street light condition weight must be much larger than the other weights in order to exhibit more resilience to Irene. In other words, the only weights that will lead to a higher level of calculated overall resilience to Irene, compared to the 2011 Snowstorm, must satisfy this relative relationship.

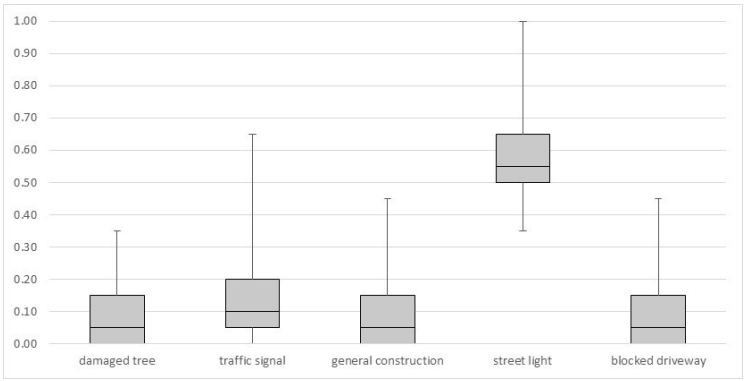

Fig. 9. Weight ranges for which there is more resilience to Irene than to the 2011 Snowstorm

4.3. Decision making support

It is important to note that such information about the range of weights within which one disaster event is considered more resilient than another is only relevant if there is both the opportunity and the desire for a decision maker, or group of decision makers, to assign different weights to the different dimensions of resilience. One of the primary advantages of creating such a differential weighting is that it allows the decision makers to reflect the relative importance of each dimension to them, in terms of the perceived contribution that it makes to the overall assessment of resilience. Because such differences may also inform the priorities of local representatives in different areas of a municipality, there may be a need for multiple decision makers to come to some sort of consensus about how much relative weight should be given to each dimension, in order to appropriately allocate resources either for mitigation in advance of a disaster event, or for recovery after the event has occurred. Using a technique such as the Analytical Hierarchy Process (AHP) can allow such a group to determine an appropriate weight value for each such dimension, and these values can then be used to generate the weighted overall resilience value along with the corresponding loss and recovery time values.

In general, however, given the complexity of assessing a multi-dimensional system, we assume that decision makers will be less interested in determining the "best" set of weights to represent their priorities, and more interested in determining if the set of weights that they actually chose is appropriately robust, subject to some uncertainty about the precision of the actual values that were determined. This implies that there is value in characterizing *changes* in the relative resilience of the different systems being compared, given adjustments in the actual values of the weights within a specified interval. The framework introduced above provides a strong foundation for doing this in a very straightforward way. We illustrate this by expanding on the previous example to include all seven disaster events being considered. Figure 10 presents the corresponding complete set of convex hulls.

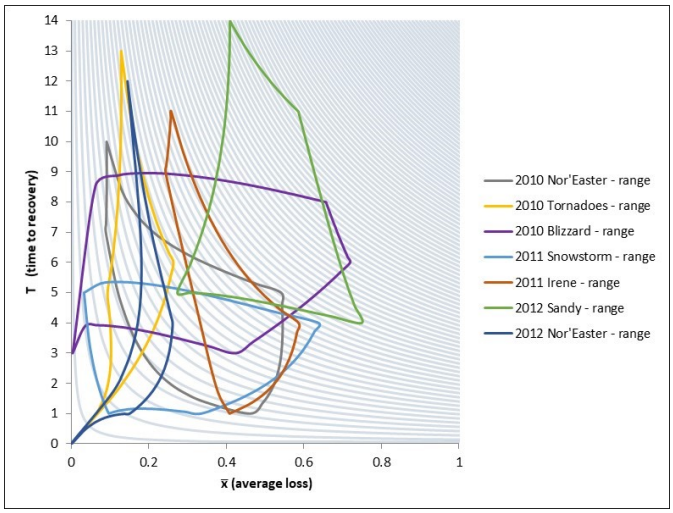

Fig. 10. Range of possible weighted resilience values for each of the seven disaster events

To illustrate the ability to characterize the impacts of adjusting weights, let us assume, without loss of generality, that an initial weighting scheme was chosen that weights each dimension equally across all seven events. The resulting relative weighted resilience values, along with their corresponding weighted average loss values and recovery times are given in Table 4.

Because we know from the example above that different weighting schemes can lead to different relative rankings for the resilience values, we may explore the robustness of these relative rankings by systematically varying the weights around their initial values. To illustrate this, we specify a range of up to +/- 50% of the original weight value within which to vary each weight, subject to the constraint that the sum of all the weight values must always equal one. Figure 11 provides a plot of the resulting set of convex combinations of the dimensions for each event.

Table 4. Overall resilience results for equal weighting scheme

Event ID	Event description	Resilience (rank)	Average loss (rank)	Recovery time (rank)
Event 1	2010 Nor'Easter	0.9099 (4)	0.2720 (4)	4.6362 (4)
Event 2	Brooklyn / Queens tornadoes	0.9458 (2)	0.1711 (1)	4.4382 (3)
Event 3	N. American Blizzard	0.7833 (6)	0.4630 (6)	6.5528 (6)
Event 4	Hurricane Irene	0.8694 (5)	0.3912 (5)	4.6725 (5)
Event 5	Major Snowstorm	0.9490 (1)	0.2347 (3)	3.0425 (1)
Event 6	Hurricane Sandy	0.7077 (7)	0.4947 (7)	8.2718 (7)
Event 7	2012 Nor'Easter	0.9372 (3)	0.2170 (2)	4.0506 (2)

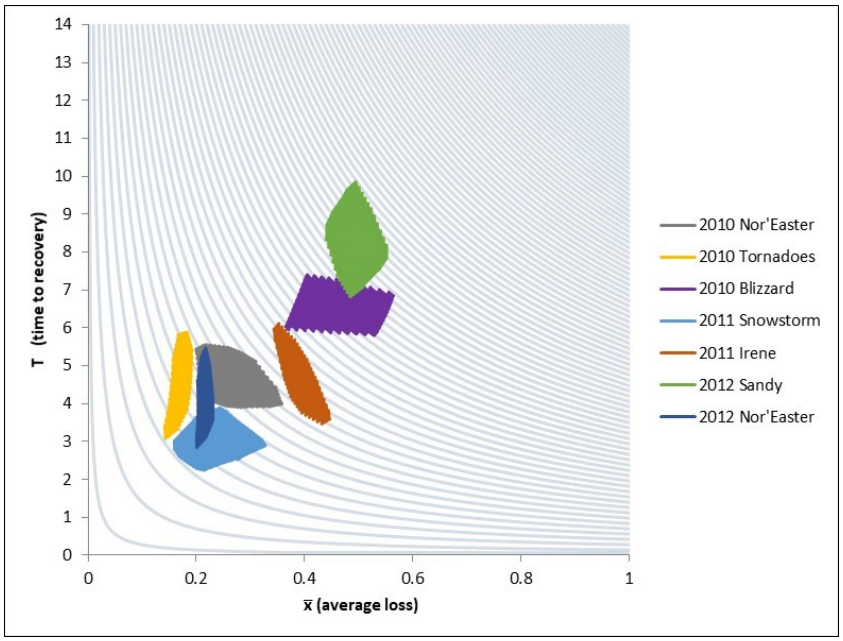

Fig. 11. +/- 50% variation weight combinations, relative to equal weighting scheme

Although Figure 11 makes it relatively straightforward to perform a high-level comparison of the ranges for the average loss and recovery time values for different events, it can also help to look at the relative distributions of those values since the resilience curves are nonlinear in nature. A graph such as Figure 12 provides the empirical distributions of values across the different ranges of weights, and better illustrates the extent to which the different events' associated resilience values overlap.

Such overlap does not necessarily indicate, however, how often the city is more resilient to one of the event than it is to another, since the different dimensions may not align with one another, as in the case

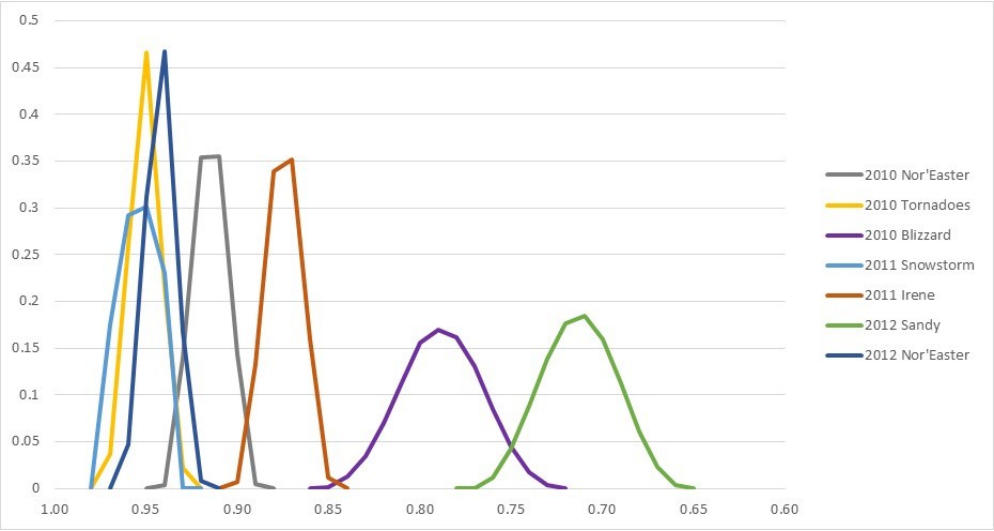

Fig. 12. Event resilience distribution

Table 5. % of weighting schemes with more resilience to row event than to column event

	2010 N'Easter	2010 Tornadoes	2010 Blizzard	2011 Irene	2011 Snowstorm	2012 Sandy	2012 N'Easter
2010 N'Easter		0.0000	1.0000	1.0000	0.0007	1.0000	0.0000
2010 Tornadoes	1.0000		1.0000	1.0000	0.4305	1.0000	0.8403
2010 Blizzard	0.0000	0.0000		0.0000	0.0000	0.9953	0.0000
2011 Irene	0.0000	0.0000	1.0000		0.0000	1.0000	0.0000
2011 Snowstorm	0.9993	0.5695	1.0000	1.0000		1.0000	0.7446
2012 Sandy	0.0000	0.0000	0.0047	0.0000	0.0000		0.0000
2012 N'Easter	1.0000	0.1598	1.0000	1.0000	0.2554	1.0000	

of the 2011 Snowstorm and Hurricane Irene that was discussed above. To get a more precise indication of the relative resilience behavior associated with the different events, therefore, we use the results of generating the range of weight values to empirically calculate the percentage of time that the resilience to each event is more than that to each of the other events, as in Table 5. These results can be consolidated into a relative ranking, as before, by simply averaging the percentage values for each row. The results of doing so, in this case, match the reported rankings that were given in Table 4.

From a decision making standpoint, Table 5 indicates the robustness of the baseline weighting scheme to changes in the relative preference for each of the weights. Thus, for example, the resilience of New York City to Hurricane Sandy was less than its resilience to any other event across all possible weight combinations that were considered, except for a very small number of cases (weight combinations) when it might be considered to be more resilient to Sandy than to the 2010 Blizzard. If we examine the actual set of weights for which this is true, it becomes clear that this behavior is exhibited only for very high weights on general construction and street light and very low weights on damaged tree and traffic signal (See Figure 13). If such a weighting scheme is not considered to be representative of the decision makers' preferences, then Sandy can definitively be considered to be the event to which the city is least resilient overall.

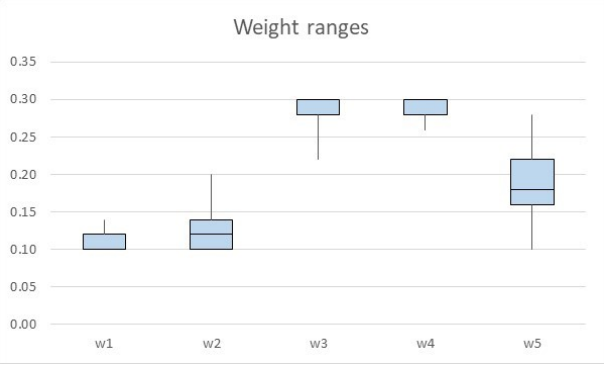

Fig. 13. Weight ranges for which NYC is more resilient to Sandy than to the 2010 Blizzard

Table 6 provides additional information to complement the relative resilience comparisons in Table 5. It gives the percentage of time that the angle of ascension is greater for the row event than for the column event, thus higher values indicate behavior skewed towards relatively longer recovery times and smaller average levels of impact. The results imply that, relative to the other events, NYC suffered more average loss and a shorter recovery time for Irene, whereas it had a longer recovery time but less average loss overall for the Tornadoes.

Table 6. Percent of weighting schemes with a larger ratio of T to X for row event than column event

	2010 N'Easter	2010 Tornadoes	2010 Blizzard	2011 Irene	2011 Snowstorm	2012 Sandy	2012 N'Easter
2010 N'Easter		0.0000	0.8909	1.0000	1.0000	0.5353	0.2562
2010 Tornadoes	1.0000		1.0000	1.0000	1.0000	1.0000	1.0000
2010 Blizzard	0.1091	0.0000		0.8897	0.8401	0.0399	0.0000
2011 Irene	0.0000	0.0000	0.1103		0.1340	0.0203	0.0000
2011 Snowstorm	0.0000	0.0000	0.1599	0.8660		0.0541	0.0000
2012 Sandy	0.4647	0.0000	0.9601	0.9797	0.9459		0.1481
2012 N'Easter	0.7438	0.0000	1.0000	1.0000	1.0000	0.8520	

Table 7. Percent of time that row has a smaller average loss than column

	2010 N'Easter	2010 Tornadoes	2010 Blizzard	2011 Irene	2011 Snowstorm	2012 Sandy	2012 N'Easter
2010 N'Easter		0.0000	1.0000	1.0000	0.1111	1.0000	0.0008
2010 Tornadoes	1.0000		1.0000	1.0000	0.9888	1.0000	1.0000
2010 Blizzard	0.0000	0.0000		0.0089	0.0000	0.7257	0.0000
2011 Irene	0.0000	0.0000	0.9910		0.0000	1.0000	0.0000
2011 Snowstorm	0.8888	0.0113	1.0000	1.0000		1.0000	0.3124
2012 Sandy	0.0000	0.0000	0.2743	0.0000	0.0000		0.0000
2012 N'Easter	0.9992	0.0000	1.0000	1.0000	0.6877	1.0000	

Table 8. Percent of time that row has a shorter recovery time than column

	2010 N'Easter	2010 Tornadoes	2010 Blizzard	2011 Irene	2011 Snowstorm	2012 Sandy	2012 N'Easter
2010 N'Easter		0.3413	1.0000	0.5415	0.0000	1.0000	0.0255
2010 Tornadoes	0.6587		1.0000	0.6963	0.0023	1.0000	0.2593
2010 Blizzard	0.0000	0.0000		0.0000	0.0000	1.0000	0.0000
2011 Irene	0.4585	0.3037	1.0000		0.0000	1.0000	0.1344
2011 Snowstorm	1.0000	0.9977	1.0000	1.0000		1.0000	0.9590
2012 Sandy	0.0000	0.0000	0.0000	0.0000	0.0000		0.0000
2012 N'Easter	0.9746	0.7407	1.0000	0.8656	0.0410	1.0000	

We may also characterize such behaviors, however, by comparing the actual loss and recovery time values across events (Tables 7 and 8). This shows that NYC had both the largest average loss values and the longest associated recovery times for Hurricane Sandy and the 2010 Blizzard, while it had the shortest recovery times overall for the 2011 Snowstorm. Table 8 further shows that the Tornadoes results were more mixed on recovery time, and thus we may conclude that it was the very small average loss values that contributed most to the ratio results in Table 6.

Tables 7 and 8 also show that the instances in which the city was less resilient to the Blizzard than to Hurricane Sandy in Table 5 must have corresponded to more weighted average loss for the Blizzard than for Sandy, since it never had a longer recovery time for the Blizzard under any of the weight combinations. Furthermore, even though there were weight combinations for which the city had less average loss for the 2010 Nor'Easter than for the 2012 Nor'Easter, and there were also weight combinations for which it had a shorter recovery time for the 2010 event than for the 2012 event, these two sets of weights never coincide. This can be seen from the results in Table 5 because NYC demonstrated less resilience to the 2010 Nor'Easter than it did to the 2012 Nor'Easter under every possible weight combination.

4.4. Additional applications of the approach

This general decision making approach can be applied in a variety of situations, depending on the focus of the decision maker or decision makers. For example, the analysis could be applied to assessing the relative resilience of different geographical or socio-political regions, such as New York City's five different boroughs, in order to compare their individual behaviors in response to the different disasters. Because even a large-scale disaster may have very different types of impacts on different areas, this can help to create a more precise view of how the community was actually affected overall.

Rather than weighting the different resilience dimensions for each event, one could instead weight the different *events* in order to determine the range of possible resilience values for each individual dimension. Such weights could be based on the relative importance of each event to the municipality, perhaps using a high-level measure such as economic damage or casualties, and they could help to uncover characteristics of the individual resilience dimensions in order to further help understand the system's overall behavior. We illustrate the potential of this in the context of the broader 311 example developed above.

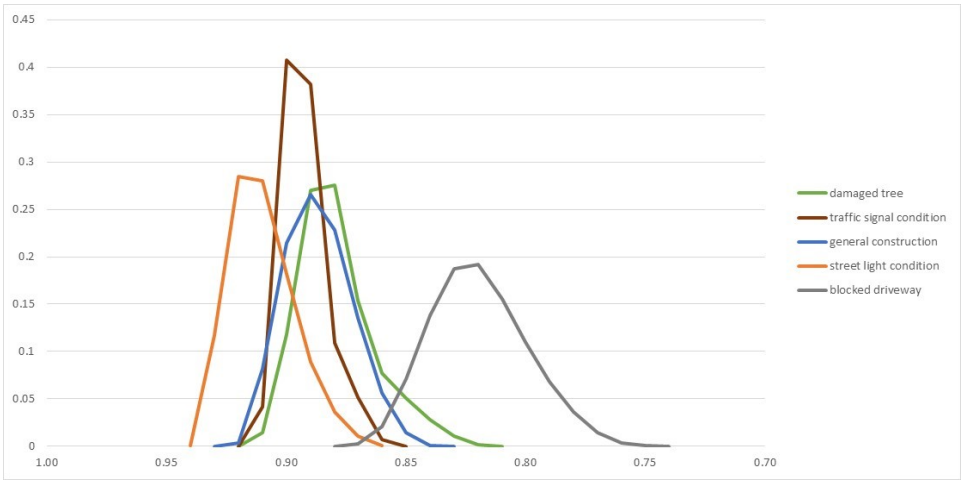

Fig. 14. Distribution of resilience values across dimensions

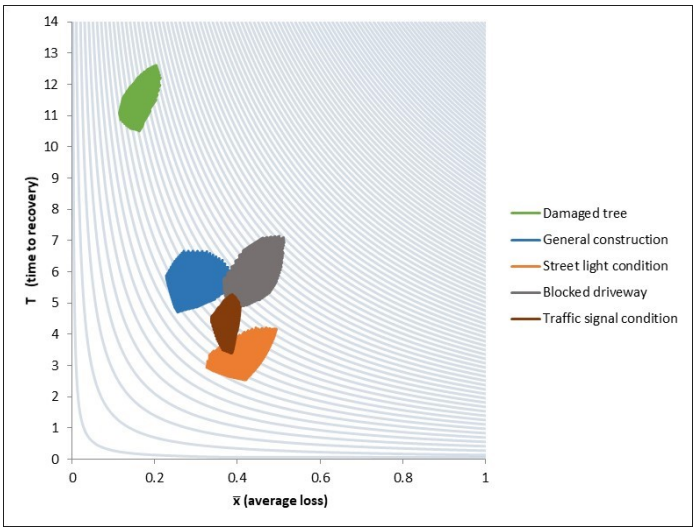

Fig. 15. +/- 50% variation weight combinations, relative to equal weighting scheme

Figures 14 and 15 give the result of varying the weights on the seven different disaster events (by +/- 50% of the equal weight values), for each of the resilience dimensions discussed above: *Damaged Tree*, *Street Light*, *Traffic Signal*, *Blocked Driveway*, and *General Construction*. From Figure 14, it can be seen that there is a lot of overlap in the distribution of resilience values across the different measures. NYC appears to have the least resilience in the *blocked driveway* dimension, in general, and the most resilience in the *street light* dimension, with the resilience values associated with *general construction* and *damaged tree*, in particular, appearing to overlap to a significant extent. It is significant, however, not only that the actual overlap between the *general construction* and *damaged tree* dimensions is only about 30%, but also that the nature of the resilience behaviors that they represent are actually distinctly different. As Figure 15 clearly illustrates, the *damaged tree* measure is associated with both much less average loss and

significantly longer recovery times than any of the other measures, despite the similarity in the actual resilience values. This clearly illustrates the value of considering not just the resilience value itself, but also the loss and recovery time from which it is constructed, as a means of better understanding how a system and its different dimensions behave in the face of complex disaster situations.

5. Discussion and conclusions

Resilience is a multi-dimensional concept that can be used to characterize the effects of a natural disaster on different complex human systems. This study provides a new, more holistic approach for explicitly capturing and quantifying the multi-dimensional nature of resilience in a rich and meaningful way, so that the relative response of different systems to a disaster or disruption can be compared much more effectively. The new approach expands on the idea of characterizing the tradeoffs between two of the main characteristics of resilient systems: the ability to resist loss and the ability to recover quickly. By considering these tradeoffs across multiple dimensions simultaneously, the approach enables decision makers to evaluate and compare the resiliency of complex systems from a number of different perspectives, and thus to form a much more in-depth understanding of those systems' responses to disaster events. This, in turn, allows for developing much more effective plans for mitigating against and responding to future disasters, since resources can be targeted towards addressing specific types of vulnerability within different system components.

One of the most significant characteristics of the approach introduced above is its support for combining the different dimensions of resilience into a single overall resilience measure, while still maintaining the ability to visualize and analytically characterize the relative behaviors of the different individual dimensions. In particular, it allows the contributions of those different dimensions to be weighted differentially, in order to reflect the relative importance of each one to the overall resilience behavior of the system. Furthermore, it supports analyzing the effects of changes in these relative weights, not just on the measure of overall resilience but also on the very different ways in which the system can resist against and then recover from a disaster event.

To illustrate the new approach, we examined the multi-dimensional resilience of New York City, in the context of seven major natural disaster events that the city experienced between 2010 and 2012. We used the number of non-emergency public service requests of different types received by 311 centers during the disaster events as a new source of data to represent the multi-dimensional impacts of those disasters on the city. The results of applying the case study clearly demonstrate the ability of the new approach to support an in-depth analysis of these different impacts, both individually and as a whole. For example, the results indicate that although the number of requests associated with damage to trees were generally less than what was observed in other dimensions, as a result of the different disaster impacts, those requests did

persist for a significantly greater length of time in every case. Such results provide important insights that can directly help with long-term planning and policy improvements, as well as with short-term resource allocations.

There are a number of different ways in which this new multi-dimensional resilience approach can be applied, beyond just case studies similar to the one that was presented. For example, although the number of service requests over time provides useful information about the current status of a community, it does not also provide direct information about how the city actually responded to the changes in the number of service requests. One interesting alternative measure to consider, therefore, is the number of service requests in process each day (Zobel et al., 2017). Such a measure could be used to assess the relative organizational resilience of the city itself, and its various agencies, in addition to that of the people in the community. There are also other publicly-generated data sets, such as message-board postings or Twitter feeds, which can provide opportunities to capture the relevant interactions between a community and the service providers that enable it to function appropriately.

More importantly, the new approach will generally apply to any complex system whose performance, and that of its relevant components (or dimensions), can be represented by time series that capture the impact of a disruption on those aspects of the system over time. The system could be a single complex entity, such as a bridge or a building that is susceptible to damage from a natural disaster, or it could be a more complex system of entities, such as a supply chain or a communication network. Our case study illustrates the use of different types of dimensional indicators that satisfy this requirement. It allows us to demonstrate that one can capture aspects of both physical characteristics (street lights, traffic signals, and general construction) and more social characteristics (damaged trees and blocked driveways as indicators of social well-being) within a complex socio-technical system.

The socially-focused indicators often adopted by other approaches to measuring multi-dimensional resilience tend to be based on relatively static data, such as that collected every 10 years by the U.S. Census Bureau [11]. In such situations, there may not be enough data points available to create a time series representation of actual loss or recovery time, even if a particular event, or time frame, is specified as the basis for measuring resilience. Acknowledging this, Zobel (2011) suggests using an entire resilience curve, rather than just a single ($\bar{X}$, T) observation, to represent the resilience behavior in a capacity-oriented context, in order to provide a means for comparing it against more output-oriented results. The newer approach developed above provides a potential opportunity means for expanding initial work such as this, and thus establishing a stronger link between these different perspectives on resilience.

One of the limitations of this new approach to characterizing multi-dimensional resilience is its inability to represent the exogenous impact of the disruptive event itself. Only the actual effect of the event on the human system is captured by the predicted resilience measure. The new approach also does not

capture potential dependencies between the amount of loss and the length of the recovery time, or the inherent dependencies between the resilience behaviors of different dimensions. In effect, it assumes that the different dimensions are independent of one another in how they respond to and then recover from a disaster. This implies that there is opportunity to extend the approach to include consideration of such interdependencies. Such an extension might allow the approach to more accurately represent cascading behaviors within systems such as supply networks, communication networks, or critical infrastructure networks, which have direct relationships between their various components. This would support more detailed analysis and comparison of different resilience policies and procedures aimed at improving the performance of such systems.

Acknowledgements

This work was supported by the National Science Foundation (NSF-CRISP #1541155 and NSF-NRT #1735139).

References

[1] UNISDR. The human cost of weather-related disasters, 1995–2015. 2015.

[2] Cutter SL, Emrich C. Are natural hazards and disaster losses in the U.S. increasing? Eos, Trans Am Geophys Union 2005;86:381. doi:10.1029/2005EO410001.

[3] Urrea G, Villa S, Gonçalves P. Exploratory analyses of relief and development operations using social networks. Socioecon Plann Sci 2016;56:27–39. doi:https://doi.org/10.1016/j.seps.2016.05.001.

[4] Gilbert S. Disaster resilience: A guide to the literature. NIST Spec Publ 2010.

[5] Boin A, McConnell A. Preparing for critical infrastructure breakdowns: The limits of crisis management and the need for resilience. J Contingencies Cris Manag 2007;15:50–9. doi:10.1111/j.1468-5973.2007.00504.x.

[6] Ouyang M, Dueñas-Osorio L. Multi-dimensional hurricane resilience assessment of electric power systems. Struct Saf 2014;48:15–24. doi:10.1016/j.strusafe.2014.01.001.

[7] Christopher M, Peck H. Building the resilient supply chain. Int J Logist Manag 2004;15:1–14.

[8] Tomlin B. On the value of mitigation and contingency strategies for managing supply chain disruption risks. Manage Sci 2006;52:639–57. doi:10.1287/mnsc.1060.0515.

[9] Li B, Hernandez I, Milburn AB, Ramirez-Marquez JE. Integrating uncertain user-generated demand data when locating facilities for disaster response commodity distribution. Socioecon Plann Sci 2018;62:84–103. doi:https://doi.org/10.1016/j.seps.2017.09.003.

[10] Hasani A, Mokhtari H. Redesign strategies of a comprehensive robust relief network for disaster

management. Socioecon Plann Sci 2018. doi:https://doi.org/10.1016/j.seps.2018.01.003.

[11] Cutter SL, Burton CG, Emrich CT. Disaster Resilience Indicators for Benchmarking Baseline Conditions. J Homel Secur Emerg Manag 2010;7. doi:10.2202/1547-7355.1732.

[12] Stewart GT, Kolluru R, Smith M. Leveraging public-private partnerships to improve community resilience in times of disaster. Int J Phys Distrib Logist Manag 2009;39:343–64. doi:10.1108/09600030910973724.

[13] UNISDR. Making Cities Resilient: Summary for Policymakers. 2013.

[14] DHS. Regional Resiliency Assessment 2009. https://www.dhs.gov/regional-resiliency-assessment-program (accessed November 8, 2017).

[15] National Research Council. Disaster resilience: a national imperative. Washington, DC: The National Academies Press; 2012.

[16] UNISDR. Terminology - UNISDR. UNISDR 2017. http://www.unisdr.org/we/inform/terminology (accessed November 2, 2017).

[17] Hosseini S, Barker K, Ramirez-Marquez JE. A review of definitions and measures of system resilience. Reliab Eng Syst Saf 2016;145:47–61. doi:10.1016/j.ress.2015.08.006.

[18] Morimoto R. Estimating the benefits of effectively and proactively maintaining infrastructure with the innovative Smart Infrastructure sensor system. Socioecon Plann Sci 2010;44:247–57. doi:https://doi.org/10.1016/j.seps.2010.07.005.

[19] Bruneau M, Chang S, Eguchi R, Lee GC, O'Rourke TD, Reinhorn AM, et al. A framework to quantitatively assess and enhance the seismic resilience of communities. Earthq Spectra 2003;19:733–52.

[20] Cutter SL, Barnes L, Berry M, Burton C, Evans E, Tate E, et al. A place-based model for understanding community resilience to natural disasters. Glob Environ Chang 2008;18:598–606. doi:10.1016/j.gloenvcha.2008.07.013.

[21] Zobel CW. Representing perceived tradeoffs in defining disaster resilience. Decis Support Syst 2011;50:394–403.

[22] Zobel CW, Khansa L. Quantifying Cyberinfrastructure Resilience against Multi-Event Attacks. Decis Sci 2012;43:687–710. doi:10.1111/j.1540-5915.2012.00364.x.

[23] Zobel CWCW, Baghersad M, Zhang Y. Calling 311: evaluating the performance of municipal services after disasters. 14th Int. ISCRAM Conf., vol. 2017–May, 2017, p. 164–72.

[24] Cutter SL. The landscape of disaster resilience indicators in the USA. Nat Hazards 2016;80:741–58. doi:10.1007/s11069-015-1993-2.

[25] Paton D, Johnston D. Disaster resilience: an integrated approach. Charles C Thomas Publisher; 2017.

[26] Vugrin ED, Warren DE, Ehlen MA. A resilience assessment framework for infrastructure and economic systems: Quantitative and qualitative resilience analysis of petrochemical supply chains to a hurricane. Process Saf Prog 2011;30:280–90. doi:10.1002/prs.10437.

[27] Nguyen K V., James H. Measuring Household Resilience to Floods: a Case Study in the Vietnamese Mekong River Delta. Ecol Soc 2013;18:13. doi:10.5751/ES-05427-180313.

[28] Zhong S, Hou X-Y, Clark M, Zang Y-L, Wang L, Xu L-Z, et al. Disaster resilience in tertiary hospitals: a cross-sectional survey in Shandong Province, China. BMC Health Serv Res 2014;14:135. doi:10.1186/1472-6963-14-135.

[29] Ambulkar S, Blackhurst J, Grawe S. Firm's resilience to supply chain disruptions: Scale development and empirical examination. J Oper Manag 2015;33:111–22. doi:10.1016/j.jom.2014.11.002.

[30] Pettit TJ, Croxton KL, Fiksel J. Ensuring Supply Chain Resilience: Development and Implementation of an Assessment Tool. J Bus Logist 2013;34:46–76. doi:10.1111/jbl.12009.

[31] Shenesey JW, Langhinrichsen-Rohling J. Perceived resilience: Examining impacts of the deepwater horizon oil spill one-year post-spill. Psychol Trauma Theory, Res Pract Policy 2015;7:252–8. doi:10.1037/a0035182.

[32] Sempier TT, Swann DL, Emmer R, Sempier SH, Schneider M. Coastal community resilience index: A community self-assessment. Mississippi. The University of Southern Mississippi and Mississippi-Alabama Sea Grant Consortium.[2015-03-24]. http://www. masgc. org; 2010.

[33] Mostofi Camare H, Lane DE. Adaptation analysis for environmental change in coastal communities. Socioecon Plann Sci 2015;51:34–45. doi:https://doi.org/10.1016/j.seps.2015.06.003.

[34] Foster KA. In search of regional resilience. Build. Reg. Resil. Urban Reg. policy its Eff., Brookings Institution Press, Washington DC; 2012.

[35] Zobel CWCW. Representing the Multi-Dimensional Nature of Disaster Resilience. 8th Int. ISCRAM Conf., 2011, p. 1–5.

[36] Zobel CW. Comparative Visualization of Predicted Disaster Resilience. Seventh Int. ISCRAM Conf., Seattle, WA: 2010, p. 1–6.

[37] Zobel CW, Khansa L. Characterizing multi-event disaster resilience. Comput Oper Res 2014;42:83–94. doi:10.1016/j.cor.2011.09.024.

[38] Chang SE, Shinozuka M. Measuring Improvements in the Disaster Resilience of Communities. Earthq Spectra 2004;20:739–55. doi:10.1193/1.1775796.

[39] Cimellaro GP, Reinhorn AM, Bruneau M. Framework for analytical quantification of disaster resilience. Eng Struct 2010;32:3639–49. doi:10.1016/j.engstruct.2010.08.008.

[40] Francis R, Bekera B. A metric and frameworks for resilience analysis of engineered and

infrastructure systems. Reliab Eng Syst Saf 2014;121:90–103. doi:10.1016/j.ress.2013.07.004.

[41] Pant R, Barker K, Zobel CW. Static and dynamic metrics of economic resilience for interdependent infrastructure and industry sectors. Reliab Eng Syst Saf 2014;125:92–102. doi:10.1016/j.ress.2013.09.007.

[42] Ouyang M, Dueñas-Osorio L, Min X. A three-stage resilience analysis framework for urban infrastructure systems. Struct Saf 2012;36–37:23–31. doi:10.1016/j.strusafe.2011.12.004.

[43] Henry D, Ramirez-Marquez JE. Generic metrics and quantitative approaches for system resilience as a function of time. Reliab Eng Syst Saf 2012;99:114–22. doi:10.1016/j.ress.2011.09.002.

[44] Adams TM, Bekkem KR, Toledo-Durán EJ. Freight resilience measures. J Transp Eng 2012;138:1403–9.

[45] Adjetey-Bahun K, Birregah B, Châtelet E, Planchet JL. A model to quantify the resilience of mass railway transportation systems. Reliab Eng Syst Saf 2016;153:1–14. doi:10.1016/j.ress.2016.03.015.

[46] Sahebjamnia N, Torabi SA, Mansouri SA. Integrated business continuity and disaster recovery planning: Towards organizational resilience. Eur J Oper Res 2015;242:261–73. doi:10.1016/j.ejor.2014.09.055.

[47] Torabi SA, Baghersad M, Mansouri SA. Resilient supplier selection and order allocation under operational and disruption risks. Transp Res Part E Logist Transp Rev 2015;79:22–48. doi:10.1016/j.tre.2015.03.005.

[48] Spiegler VLM, Naim MM, Wikner J. A control engineering approach to the assessment of supply chain resilience. Int J Prod Res 2012;50:6162–87. doi:10.1080/00207543.2012.710764.

[49] Zobel CW. Quantitatively Representing Nonlinear Disaster Recovery. Decis Sci 2014;45:1053–82. doi:10.1111/deci.12103.

[50] Zobel CW, Baghersad M, Zhang Y. An approach for quantifying the multidimensional nature of disaster resilience in the context of municipal service provision. Urban Disaster Resil. Secur., 2017.

[51] City of New York. Mayor Bloomberg announces high NYC 311 customer satisfaction | City of New York 2013. http://www1.nyc.gov/office-of-the-mayor/news/382-13/mayor-bloomberg-nyc-311-customer-satisfaction-survey-shows-high-customer-satisfaction/ (accessed November 6, 2017).

[52] Meinshausen N. Quantile Regression Forests. J Mach Learn Res 2006;7:983–99.

[53] Barber CB, Dobkin DP, Huhdanpaa H. The quickhull algorithm for convex hulls. ACM Trans Math Softw 1996;22:469–83. doi:10.1145/235815.235821.