8-11-2021

Analysis of the Effects of Spatiotemporal Demand Data Aggregation Methods on Distance and Volume Errors

Zachary Hornberger

Bruce A. Cox
Air Force Institute of Technology

Raymond R. Hill
Air Force Institute of Technology

Analysis of the effects of spatiotemporal demand data aggregation methods on distance and volume errors

Zachary Hornberger, Bruce Cox and Raymond R. Hill
*Department of Operational Sciences, Air Force Institute of Technology,
Wright-Patterson AFB, Ohio, USA*

Abstract

Purpose – Large/stochastic spatiotemporal demand data sets can prove intractable for location optimization problems, motivating the need for aggregation. However, demand aggregation induces errors. Significant theoretical research has been performed related to the modifiable areal unit problem and the zone definition problem. Minimal research has been accomplished related to the specific issues inherent to spatiotemporal demand data, such as search and rescue (SAR) data. This study provides a quantitative comparison of various aggregation methodologies and their relation to distance and volume based aggregation errors.

Design/methodology/approach – This paper introduces and applies a framework for comparing both deterministic and stochastic aggregation methods using distance- and volume-based aggregation error metrics. This paper additionally applies weighted versions of these metrics to account for the reality that demand events are nonhomogeneous. These metrics are applied to a large, highly variable, spatiotemporal demand data set of SAR events in the Pacific Ocean. Comparisons using these metrics are conducted between six quadrat aggregations of varying scales and two zonal distribution models using hierarchical clustering.

Findings – As quadrat fidelity increases the distance-based aggregation error decreases, while the two deliberate zonal approaches further reduce this error while using fewer zones. However, the higher fidelity aggregations detrimentally affect volume error. Additionally, by splitting the SAR data set into training and test sets this paper shows the stochastic zonal distribution aggregation method is effective at simulating actual future demands.

Originality/value – This study indicates no singular best aggregation method exists, by quantifying trade-offs in aggregation-induced errors practitioners can utilize the method that minimizes errors most

relevant to their study. Study also quantifies the ability of a stochastic zonal distribution method to effectively simulate future demand data.

1. Introduction

Location modeling is a branch of operations research with vast real-world applicability and thus has been studied for a number of decades. Location modeling typically considers the location and time of demand signals over a network and optimizes the corresponding location of a servicing asset, such as a factory or vehicle. The underlying spatiotemporal demand signal data points are thus instrumental to the quality of the resulting model.

When considering extremely large spatiotemporal data sets, there is frequently a need to aggregate demand points to make the problem more tractable for the solver, clearer for the analyst and comprehensible for the end-user. While heuristic solvers can be used to overcome tractability issues by finding near-optimal solutions on the complete data set, they do not alleviate all needs for aggregation. Of specific interest, stochastic spatiotemporal demand data sets, such as search and rescue (SAR), medevac, ambulance or taxi data, may require the creation of aggregated zones as a first step in either a stochastic location optimization problem or a demand prediction model.

Aggregation, while of practical use, is not a lossless compression, and introduces aggregation error into the model. When location data are aggregated, the resulting grouping's location is traditionally represented by an aggregated data point. The distances between the actual demand points and the aggregated data points depend on the size of the aggregated region and the manner of aggregation. Similarly, the magnitude of uncertainty in the aggregated demand volumes is influenced by the nature of the aggregation. Therefore, special consideration must be given to the aggregation technique used when solving location problems.

The impact of aggregation becomes more pronounced when the geographic region expands in size and there is high variability in demand density across the region; this struggle is actualized in studying the United State Coast Guard (USCG) District 14's SAR mission. The international community recognizes the need for global cooperation in responding to emerging crises around the world. Nations have entered into SAR agreements, dividing the globe into respective search and rescue regions (SSRs). Per the United States National Search and Rescue Supplement to the International Aeronautical and Maritime Search and Rescue Manual (US Coast Guard, 2013), the USCG is the federal SAR coordinator for SAR missions within the United States' maritime SSRs and the aeronautical SSRs that do not overlay the continental USA or Alaska.

USCG District 14 is headquartered in Honolulu, HI and is responsible for USCG statutory missions across the Pacific region. In particular, the district's SSR spans more than 12 million square nautical miles, although the preponderance of SAR emergencies occur in the vicinity of Guam and the Hawaiian Islands. Additionally, District 14 has among the fewest assets in the USCG fleet, increasing the necessity to optimally posture those assets across the Pacific. Given the time-sensitive nature of rescue operations, it is imperative the USCG be optimally postured to ensure rapid response. Over the past decade, researchers have partnered with Coast Guard units – USCG and international – to solve these variations of the traditional facility location problem. These studies typically use historic SAR event data as the foundation of either a deterministic or simulation-based location model.

This study makes the following contributions: First, it quantifies trade-offs in both the distance and volume errors for six quadrat aggregation methods of varying fidelities as well as two intentional zonal distribution aggregation models using District 14 SAR emergency data as a practical basis for consideration. Second, it quantifies the ability of the stochastic zonal distribution method (SZDM) to effectively simulate future demand data.

The remainder of the paper proceeds as follows: Section 2 of this paper reviews previous works related to the aggregation of data for location models in general and Coast Guard SAR missions in particular. In Section 3, we outline the methodology for implementing various aggregation techniques, both deterministic and stochastic, using a training data set. In Section 4, we evaluate the effectiveness of these techniques by quantifying the aggregation errors between the modeled demand and actual demand over a two-year period. In Section 5, we review our findings and provide recommendations for future research.

2. Related works
Researchers have long been cognizant of a relationship between the methods used to aggregate location data and the resulting solutions generated by location models using this data. In the fields of geography and ecology, aggregation error of spatial data points is dubbed the modifiable areal unit problem (MAUP) (Openshaw, 1984; Dark and Bram, 2007) or the zone definition problem (Fotheringham *et al.*, 1995; Curtis and MacPherson, 1996). Research into MAUP typically decomposes the problem into two main effects: the scale effect and the zone effect. The scale effect refers to the impact on the spatial analysis results that are caused by the fidelity of the aggregation; for example, the impact of aggregating demand in a city using $200 \, \text{m} \times 200 \, \text{m}$ grids versus $1 \, \text{km} \times 1 \, \text{km}$ grids. Conversely, the zone effect refers to the impact caused by the way in which aggregation zones are bounded; for example, the impact of aggregating demand in a state using county lines versus city limits versus a grid overlay (Openshaw, 1984; Dark and Bram, 2007). Jelinski and Wu (1996) created a seminal contrived demonstration of these effects, which we replicate for completeness in Figure 1.

Previous research on MAUP has cautioned against arbitrary aggregation of spatial data and stressed its threat on the reliability of the resulting location analysis. Openshaw (1984) was foundational in the study of MAUP and called for developing better methods for aggregating spatial data owing to MAUP's impact on the reliability of geographic studies. Curtis and MacPherson (1996) studied data for New York and concluded that researchers can bias the results of their analysis based on the means of aggregation, even if there appears to be a logical basis for the used method of aggregation. Fotheringham *et al.* (1995) go so far as to question the accuracy of any location-based analysis conducted using aggregated data because of the effects of MAUP.

In studies of MAUP, and aggregation theory in general, trends have emerged. Increases in the number of aggregated zones are typically proportional to decreases in distance-based aggregation error; distance-based aggregation error disappears when each distinct demand point is assigned to a unique zone (i.e. the number of aggregation zones equals the number of demand points). As any grouping introduces an associated level of distance-based error, it follows that reducing the amount of aggregation would subsequently reduce this error. Francis *et al.* (2004) note the law of diminishing returns applies in this context, however, suggesting that iterative reductions in the number of aggregate groups show diminishing improvements to error reduction. Francis and Lowe (1992) discuss the *paradox of aggregation*, noting that solving formulations to minimize error can be more cumbersome than the original location problem being solved, which is counter-intuitive as aggregation is

used to simplify the resolution of these original location problems. Dark and Bram (2007) consider the trends corresponding to both the scale effect and the zone effect. A known benefit of aggregation is tied to the scale effect; predictions of aggregated demand levels tend to be more accurate with fewer, larger aggregate zones. This accuracy increase is when there are more demand points consolidated in each zone, the demand variance between zones decreases. The impact of zone effect is less understood and tends to differ from problem-to-problem.

The importance of careful aggregation has been thoroughly studied and is synthesized in Francis *et al.* (2009). In their survey of previous literature regarding aggregation error associated with location problems, Francis *et al.* note that there is an inherent trade-off when aggregating data points. Aggregation has the benefit of decreasing computational requirements and any confidentially concerns. Francis *et al.* also note that aggregation tends to decrease the statistical uncertainty of the data by amassing larger collections of data points, thus increasing the size of the pooled sample and subsequently reducing the sample standard deviation. Unfortunately, aggregation also increases the error within the model by introducing aggregation error. Thus, there does not exist a singular "best" level of aggregation, and the trade-offs inherent in aggregation must be considered.

In addition to the theoretical work on this problem, there has been applied work specifically relating to Coast Guard SAR missions, the preponderance of which has emerged in the past decade. Studies researching the allocation of SAR assets, or facilities, typically adopt a quadrat modeling technique for aggregating location data (Akbari *et al.*, 2018a, 2018b; Karatas *et al.*, 2017; Afshartous *et al.*, 2009). This technique consists of decomposing the region in question into square cells using a grid overlay.
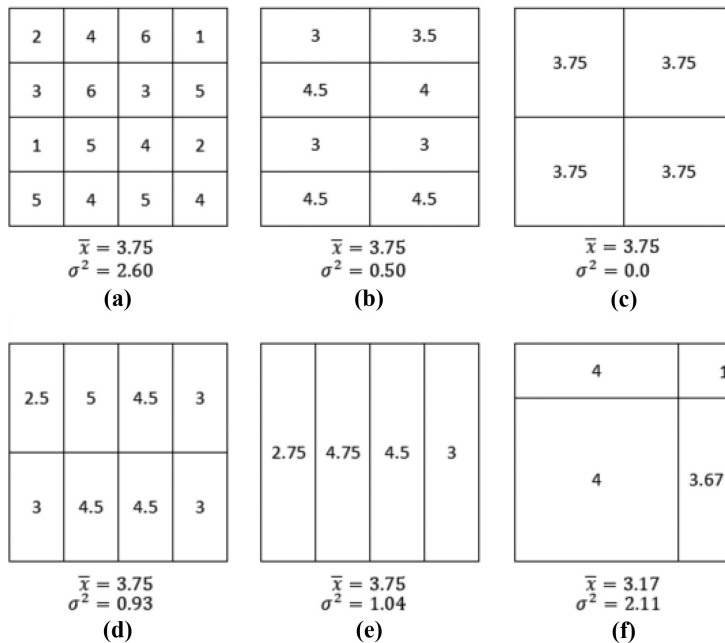

Figure 1.
(a–c) Scale effect. As scale of aggregation increases, mean does not change but variance declines. (d, e) Aggregation effect. Keeping scale equal but changing method of aggregation changes variance. (c, e, f) Even when the number of zones is constant (4) mean and variance can change

Armstrong and Cook (1979) constructed a goal-programming model for assigning SAR aircraft, incorporating probabilistic consideration for the time required by the aircraft to locate distress events in different areas of the corresponding region, using a grid overlay to create a collection of square zones. These zones were then assigned deterministic values, representing the average number of distress events per month. Similarly, Karatas *et al.* (2017) used a quadrat model for simulating the location and volume of distress calls for the Turkish Coast Guard in the Aegean Sea. They first determined the optimal resource allocation strategy using individual events as separate demand nodes and then evaluated the effectiveness of this strategy using simulated demand.

Notably, the quadrat method is frequently adopted in crime data analyses, which typically seek to quantify spatial trends in criminal activity across a city or state (Anselin *et al.*, 2000; Eck *et al.*, 2005). The incorporation of kernel density estimation with the quadrat model, popular in crime data analysis (Anselin *et al.*, 2000), has been previously implemented in SAR location problems. The kernel density estimation method composes the region into grid cells and assigns a density function to each data point (s_i). Points that are within proximity to each other relative to a specified bandwidth (τ) are grouped into a kernel (k) and their density functions are combined. The resulting image is a smooth heat map with greater densities illustrated over areas that have the most activity clustered closely together (Anselin *et al.*, 2000; Eck *et al.*, 2005). Erdemir *et al.* (2008) used kernel density estimation when considering the problem of locating aeromedical bases across the state of New Mexico. Similarly, Akbari *et al.* (2018b) implemented a kernel density estimation approach to approximate the intensity of distress calls received by the Canadian Coast Guard. They varied the size of the grid overlay based upon the proximity to the shoreline. This decision was based upon the assumption that because most distress events occurred closer to shore, the analysis would benefit from greater fidelity in aggregation along the coastline.

While most work regarding SAR posturing has incorporated quadrat techniques, Azofra *et al.* (2007) introduced an intuitive method that has been applied to maritime SAR. Instead of defaulting to grids, the Azofra *et al.* zonal distribution model (ZDM) allows for flexibility in the definition of emergency zones, such as zones based upon subject matter expertise. Once the zones are determined, the centroids of distress calls, dubbed *superaccidents*, are computed for each zone. The ZDM is a gravitational model, with the determination in optimal SAR operational response based upon the distance to the superaccidents and their associated weight. They demonstrate the implementation of this model using a notional example involving three superaccidents and three ports.

Since the introduction of the ZDM, some researchers have opted to expand upon it by applying it to real-world problems. Ai *et al.* (2015) use this model for locating supply bases and positioning vessels for maritime emergencies for a portion of the coastline of China along the Yellow Sea. While not adhering to the strict grid cells of previous studies, their zones remained rectangular in shape and varied in size across the region. Razi and Karatas (2016) improved upon the ZDM by using a k-means clustering algorithm for defining the zones and implementing a weighted approach for locating the superaccidents. By adopting this approach, Razi and Karatas define the aggregated zones and corresponding representative demand nodes based upon historical trends in distress calls in the Aegean Sea rather than arbitrary cells. Hornberger *et al.* (2021) propose an extension to the work of Razi and Karatas, which they dub the stochastic ZDM. Their model implements hierarchical k-means clustering algorithm to define the aggregation zones, fits probability distributions to model the SAR demand for each zone and then uses empirically constructed discrete distributions to model the corresponding rescue response for each emergency.

A review of the existing literature regarding SAR asset posturing models, and indeed emergency asset posturing models, finds a lack of explicit consideration regarding the impact of aggregation. Additionally, as SAR research expands to larger regions of consideration (e.g. oceans vs. seas or shorelines), it is necessary to more thoroughly consider the effects of various aggregation methods. Outside of SAR, and more generally emergency response asset modeling (Araz et al., 2007), other transportation resource posturing problems which use massive demand data sets assume or require demand aggregation (e.g. taxi service areas, Li and Szeto, 2019; Rajendran and Zack, 2019) and should also be concerned with how such aggregation effects the associated location modeling. To provide such consideration, our study uses historic SAR data from across the Pacific Ocean to compare the effectiveness of a zonal aggregation technique compared to quadrats of varying fidelity. Additionally, we evaluate these trade-offs in the aggregation as applied to deterministic and stochastic implementations.

3. Methodology
In this section, we consider two key characteristics that define a zonal aggregation of demand signals: dividing the region into zones and modeling the demand level. Using these two characteristics as the framework, we model and compare the following methodologies: deterministic quadrat approaches of various fidelities, the Razi and Karatas (2016) ZDM and the Hornberger et al. (2021) stochastic ZDM.

These methodologies are compared using the District 14 SAR region, an interesting test case owing to its large area and highly variable demand levels; Figure 2 depicts the Honolulu Maritime Search and Rescue Region (US Coast Guard, 2014). Historic SAR demand data were obtained from the Marine Information for Safety and Law Enforcement (MISLE) database to form both a training set and a test set. The training set is comprised of SAR events from a five-year span (January 2011 to December 2015) and is used to construct the models of spatiotemporal SAR demand. The accuracy of the aggregated demand methodologies is then evaluated using the test set which represents the historic SAR data for the same region from January 2016 to December 2017. A total of 500 emergency events were drawn from the test set, with replacement, and the actual number and location of events compared versus predicted. This process was repeated for 10,000 draws.

The training and test data are scoped to only consider events that occurred within the District 14 area of responsibility (AOR). Additionally, demand points missing global
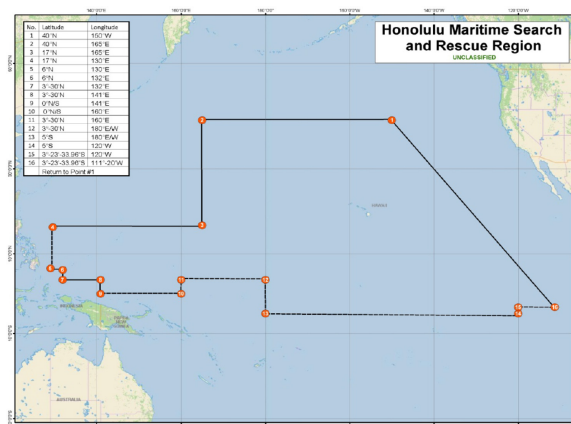
Figure 2.
Honolulu Maritime
SAR Region

positioning system coordinates were removed as were data points classified as medical consultations as these consultations only require a discussion with a medical professional over the phone and resources are not dispatched. The final training set contains 2,629 demand points, and the test set contains 1,080 demand points.

3.1 Modeling spatiotemporal demand

3.1.1 Quadrat models. The quadrat aggregation approach was implemented with six different quadrat scales to test the impact of the scale effect. These six grid-based decompositions of the region are labeled Aggregations A–F. Aggregation A considered the region of study as a singular zone, consolidating all demand points (Figure 3). Aggregation B divided the region into two zones along the antimeridiean (Figure 4). Aggregations C–E are iterative increases in fidelity, decomposing the region into 8, 15 and 43 zones, respectively (Figures 5–7). Aggregation F adopts the approach used by Akbari *et al.* (2018b) and allows for smaller grid cells in sections of higher demand. Specifically, the two zones from Aggregation E with the greatest proportion of Guam and Hawaiian Island workloads are further decomposed into $1° \times 1°$ cells; Aggregation F results in 212 zones. Aggregation F is depicted in Figure 8.

3.1.2 Zonal distribution models. Aggregation ZDM was constructed utilizing the implementation of the ZDM as established in Razi and Karatas (2016), with one exception (see below). The AOR was divided using a weighted *k*-means clustering algorithm (Figure 9). The one exception is that Razi and Karatas defined the weight of each SAR event
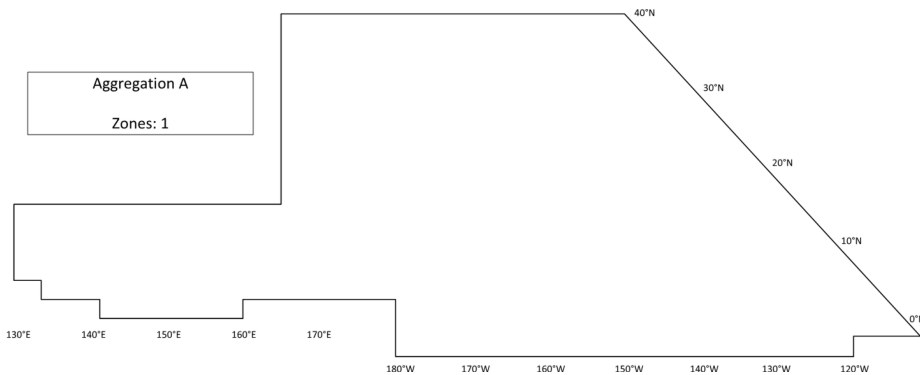

Figure 3.
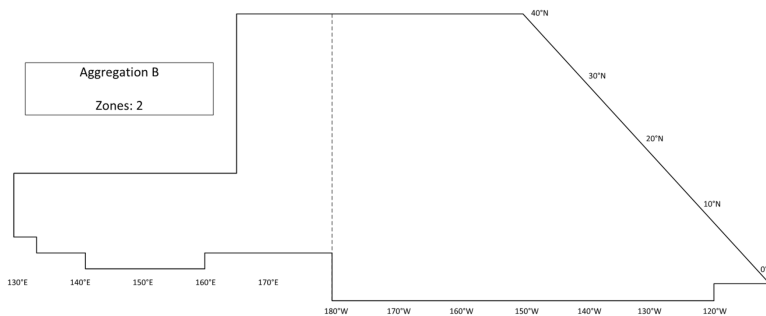Aggregation A (one zone)

Figure 4.
Aggregation B (two zones)

using an analytical hierarchy process based upon the level of fatality, material damage, response arduousness and environmental impact. Their weighting scheme was not viable for this study based on the available information in MISLE, so this implementation of Razi and Karatas's procedure uses *total activities* as a weighting metric. The metric of total activities represents the number of resources assigned to a rescue operation, in addition to the instances when a significant change occurred in the course of the rescue operation; this metric of total activities serves as a proxy for the complexity of a SAR event. Razi and
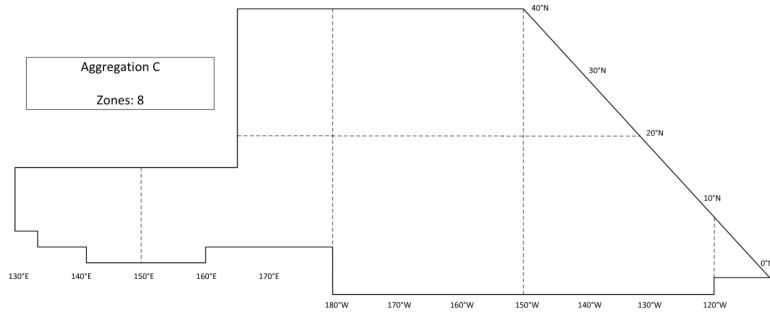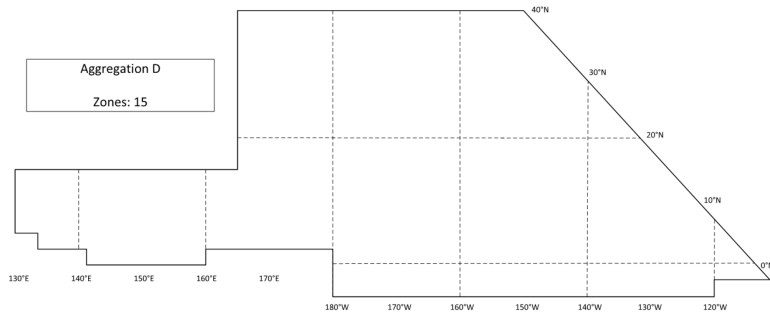
Figure 5.
Aggregation C (eight zones)

Figure 6.
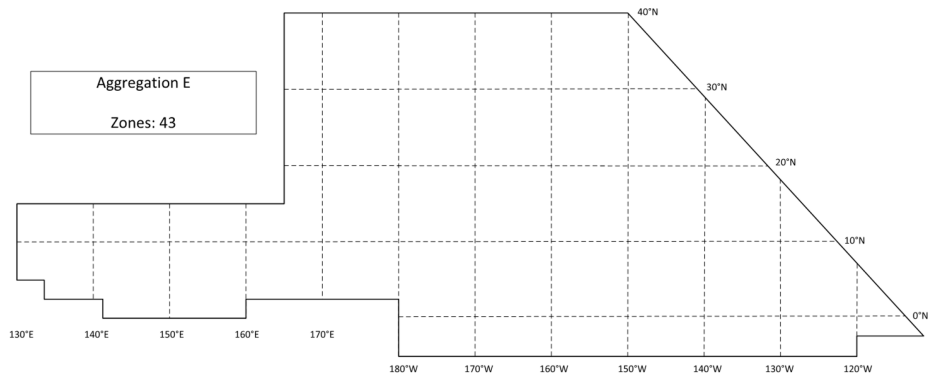Aggregation D (15 zones)

Figure 7.
Aggregation E (43 zones)

Karatas determine the number of zones to cluster demand points into based upon a *rule of thumb method* proposed by Kodinariya and Makwana (2013). This method suggests that the number of zones Z is based upon the total number of events K, such that $|Z| \approx \sqrt{|K|/2}$.

In addition to aggregation ZDM as defined above, another ZDM-inspired aggregation was created using same process as above but with only 15 zones. This aggregation is referred to as ZDM-15. A figure for ZDM-15 is not included for space considerations, but resembles that of Figure 9 with 15 versus 36 zones.

Aggregation SZDM was developed by implementing the stochastic ZDM approach proposed by Hornberger *et al.* (2021) (Figure 10). Hornberger *et al.* used a hierarchical *k*-means clustering algorithm to aggregate demand points into zones. All demand points are sorted into mutually exclusive groups based upon the unit that coordinated the response and the types of assets used in the response. District 14 is divided into Sector Guam and Sector Honolulu, which split the coverage of the AOR around longitude 160° E. Current policy dictates that the mission range for USCG boats is 50 nautical miles from the shoreline of an island on which there exists a USCG boat station; District 14 has boat stations located on the islands of Guam, O'ahu, Kaua'i and Maui. Hornberger *et al.* note that a reasonable approximation of asset utilization would be a combination of boats and helicopter aircraft responding to SAR events within the 50 nautical mile boundary of these islands while a combination of cutters and aeroplane aircraft respond to SAR events beyond these
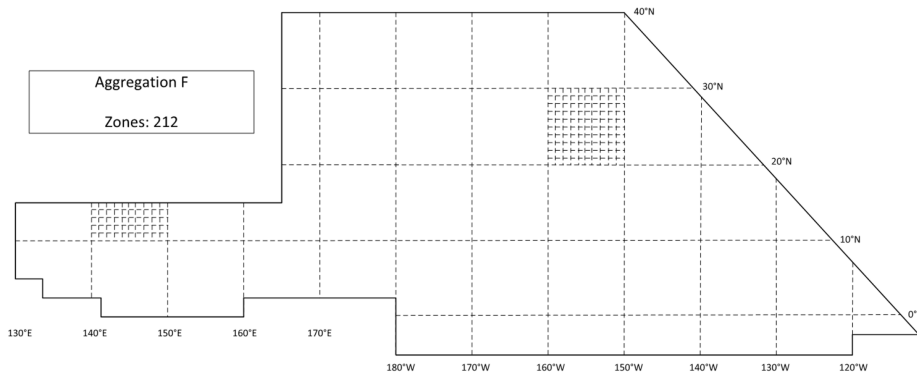

Figure 8.
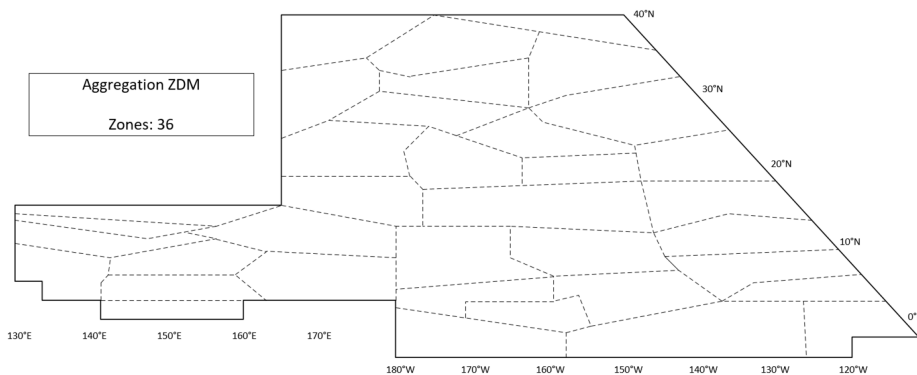Aggregation F (212 zones)

Figure 9.
Aggregation ZDM (36 zones)

boundaries. Therefore, all demand points were sorted into the following mutually exclusive groups: Guam boat/helicopter events, Guam cutter/airplane events, HI boat/helicopter events and Hawaii cutter/airplane events. These groups are further decomposed into clusters based upon the geographic proximity of the data points by using a *k*-means clustering algorithm. The number of zones was determined by considering the relationship between the number of zones and the corresponding within-cluster variance. A plot of this relationship forms an *elbow curve*, whose name is tied to the phenomena that initial groupings account for a greater reduction in variance compared to subsequent groupings, i.e. the "elbow" of the curve represents the point of diminishing returns. The "elbow" of the curve occurs at the suggested number of zones for the data set. Figure 11 shows the elbow curve for the Hawaiian cutter event cluster. The number of potential clusters is on the *x*-axis, whereas the within-cluster variance is on the *y*-axis. Note that as the number of clusters increases the within-cluster variance decreases in a convex fashion. This phenomenon enables the analyst to pick the number of clusters associated with the "crook of the elbow." In particular for this zone the authors choose six clusters, as this is the last "joint" before the curve levels off. This process, while an effective rule of thumb, is subjective.

3.2 Methods of comparative analysis

This study evaluates the effectiveness of various methods of aggregation when conducting spatiotemporal forecasting. Specifically, we assess the merit of the Razi and Karatas (2016) deterministic zonal distribution mode, and the Hornberger *et al.* (2021) stochastic ZDM, comparing their effectiveness against traditional quadrat methods of varying fidelity's. To conduct these comparisons, two metrics are considered: distance-based aggregation error and volume-based aggregation error.

The distance-based aggregation error (d_e) represents the total distance [as calculated using the Haversine formula, equation (2)] between where events were modeled as occurring ($\hat{x}_j$) and the actual location of their occurrence ($x_{i,j}$), for each event ($i \in I$) in the zone ($j \in J$). The anticipated event locations for all zones are weighted centroids for each zone. In the quadrat models, the centroids are computed as an average of the latitudes/longitudes, multiplied by the events' corresponding total activities, for all events in the zone. In the zonal and stochastic ZDM, the clustering algorithm yields a weighted centroid. The distance-based aggregation error metric is as follows:
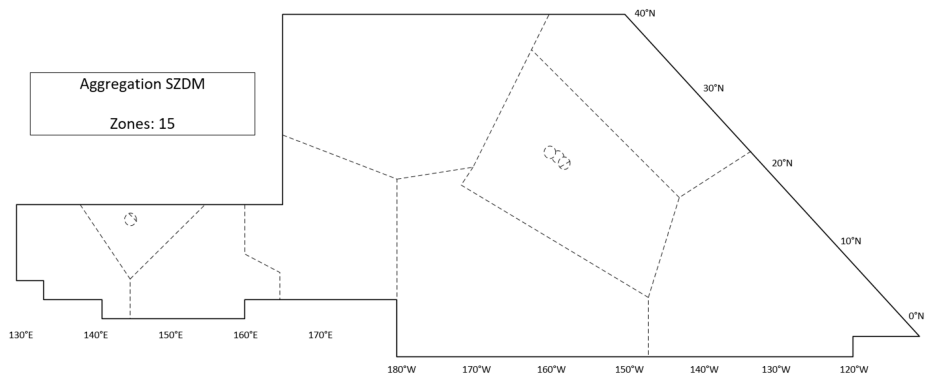

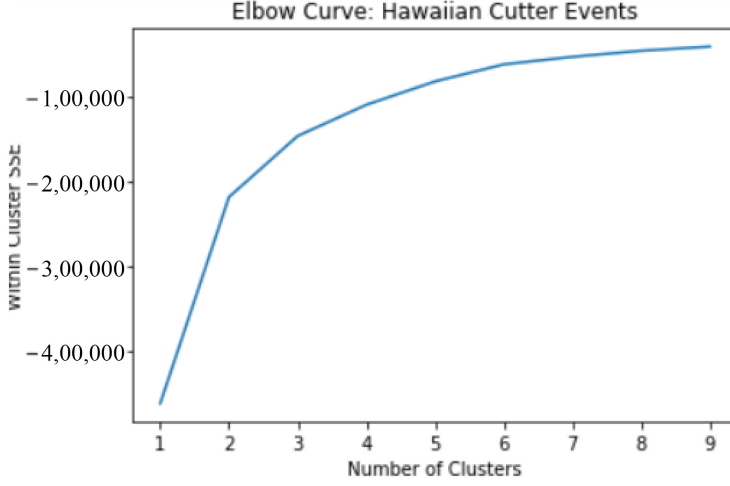**Figure 10.**
Aggregation SZDM
(15 zones)

Figure 11.
Sample "elbow" curve
for the Hawaii cutter
events cluster

$$d_e = \sum_{i \in I} \sum_{j \in J} |x_{i,j} - \hat{x}_j| \qquad (1)$$

where the Haversine formula:

$$d = 2r \arcsin\left(\sqrt{\sin^2\left(\frac{\phi_2 - \phi_1}{2}\right) + \cos(\phi_1)\cos(\phi_2)\sin^2\left(\frac{\theta_2 - \theta_1}{2}\right)}\right) \qquad (2)$$

which, given latitudes ϕ, and longitudes θ, calculates the great-circle distance between two points, is used to calculate each individual distance.

The weighted distance-based aggregation error (d_{we}) is the sum of the differences in distance between where individual assets are modeled as being deployed to ($\hat{x}_j$) and the actual location assets are dispatched to. The weighting (w_i) is the number of assets assigned to the rescue operation. The difference between d_e and d_{we} is that the former treats individual SAR events as being equal in magnitude, whereas the latter incorporates the number of deployed assets. As with d_e the individual distances in d_{we} are calculated using the Haversine formula:

$$d_{we} = \sum_{i \in I} \sum_{j \in J} w_i |x_{i,j} - \hat{x}_j| \qquad (3)$$

The distance-based aggregation error (d_e) and the weighted distance-based aggregation error (d_{we}) are both computed for all aggregations A–F, as well as for the ZDM and the SZDM.

The volume-based aggregation error (v_e) represents the total difference between the predicted level of monthly demand ($\hat{l}_{j,k}$) and the actual level of monthly demand (l_j), for each month in the considered time frame ($k \in K$). The metric is computed as follows:

$$v_e = \sum_{j \in J} \sum_{k \in K} |\hat{l}_{j,k} - l_j| \qquad (4)$$

Given that a primary difference between ZDM and SZDM is the integration of stochastic elements in the modeling of the demand, both deterministic and stochastic demand comparisons for volume-based aggregation error are conducted. For purposes of consistency, all frequency considerations are made on a *per month* basis.

Mean distance error and mean weighted distance error were calculated for all 10,000 bootstrapped samples of 500 emergency events for aggregations A–F, ZDM, ZDM-15 and SZDM. This process enables the creation of a distribution of the mean errors, and leveraging the central limit theorem 95% confidence interval distributions were created around the mean errors for all distributions. This enables both visual and statistical comparison of the aggregations.

4. Analysis
4.1 Distance-based aggregation error
The mean distances, in nautical miles, between the aggregated demand point and the subsequent demand nodes for the 10,000 samples are shown in Table 1. This table also provides upper and lower bounds for a 95% confidence interval around the mean distance error. The mean distance errors are also plotted and the resulting distributions presented in Figure 12.

The resulting distance-based aggregation error for the quadrat models (d_e) reflects the law of diminishing returns, as described by Francis *et al.* (2004). Aggregations ZDM and SZDM perform very well compared to the quadrat models. The ZDM has a lower associated location error than Aggregation E, despite only having 36 zones compared to Aggregation E's 43 zones. This runs counter to the general claim that more zones always improve the accuracy of the location model, suggesting instead that deliberate steps can be implemented to aggregate spatial demand points in fewer clusters while still achieving competitively low levels of location error. The stochastic ZDM's results support this observation, achieving a 41.7% reduction in distance-based aggregation error compared to Aggregation D despite using the same number of zones.

Similar trends are observed when the attention is shifted from the error in SAR event distances to the error in resource dispatch distances (as seen by the weighted distance error d_{we}), as seen in Table 2 and Figure 13. There is a steady improvement in accuracy as the number of zones is increased, with the exception of aggregations ZDM and SZDM. Additionally, the differences between d_e and d_{we} are notably larger for the quadrat models compared to aggregations ZDM and SZDM; the stochastic ZDM had the smallest increase in location error when weighting by the number of resources dispatched. These observations suggest that deliberate zoning of demand point can enhance the robustness of aggregate zones to weighted events, particularly when the zones are developed with consideration to

Aggregation	No. of zones	Lower bound	Mean	Upper bound
A	1	1,347.8	1,386.9	1,426.0
B	2	224.9	244.6	264.4
C	8	158.5	170.5	182.6
D	15	137.7	148.2	158.7
E	43	72.4	77.1	81.9
F	212	49.8	54.6	59.3
ZDM	36	49.8	54.8	59.6
ZDM-15	15	86.7	93.6	100.5
SZDM	15	80.9	90.0	99.1

Table 1.
Mean distance-based aggregation error for 10,000 samples of size 500. Lower and upper bounds for 95% confidence interval and mean are provided
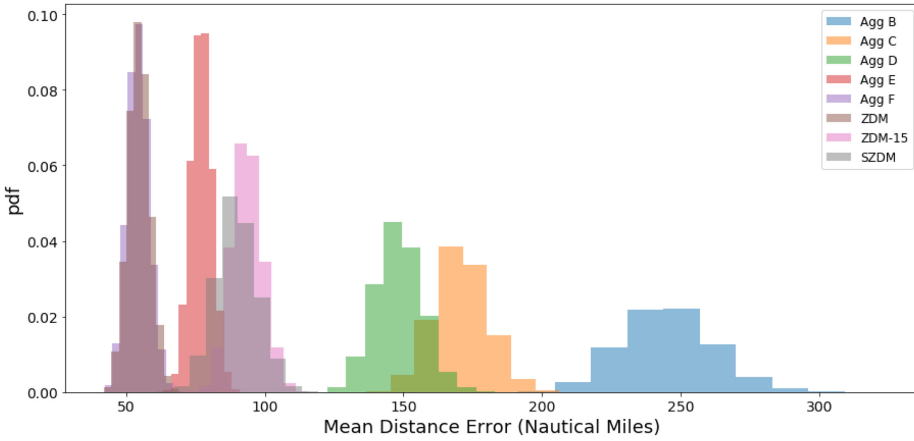
Figure 12.
Distribution of mean errors for all aggregations

both geographic proximity and the operational characteristics that are tied to the event weights.

4.2 Deterministic volume-based aggregation error

A comparison of deterministic demand models was conducted using both the mean and median monthly demand volumes for each zone in Aggregations A–F and ZDM. The average monthly error in volume based upon mean and median monthly demand for each zone compared to the actual demand volumes is depicted in Table 3. The phenomena described by Francis *et al.* (2009) and Dark and Bram (2007) is observed; there is a general increase in total volume-based aggregation error as the number of zones increases.

Consistent across all aggregation methods is the lower demand error when defining the anticipated volume of cases using the median instead of the mean. This trend becomes more pronounced as the demand error increases, suggesting that greater improvements to deterministic demand modeling can be achieved using the median when the number of zones within an aggregation models are higher versus aggregation models with fewer divisions. The advantage of using the median zonal demand versus the mean zonal demand appears to derive from the resistance of the median to extreme events. Whereas the mean demand level is impacted by exceptionally busy or slow months of SAR activity, the median provides a more stable metric.

Aggregation	No. of zones	Lower bound	Mean	Upper bound
A	1	3,416.4	3,709.3	4,002.3
B	2	515.0	577.0	639.0
C	8	377.5	422.7	468.0
D	15	326.9	370.7	414.5
E	43	177.2	197.1	217.0
F	212	119.3	138.3	157.3
ZDM	36	111.9	125.3	138.7
ZDM-15	15	198.3	217.1	235.9
SZDM	15	177.3	200.3	223.4

Table 2.
Weighted distance-based aggregation error for 10,000 samples of size 500. Lower and upper bounds for 95% confidence interval and mean are provided

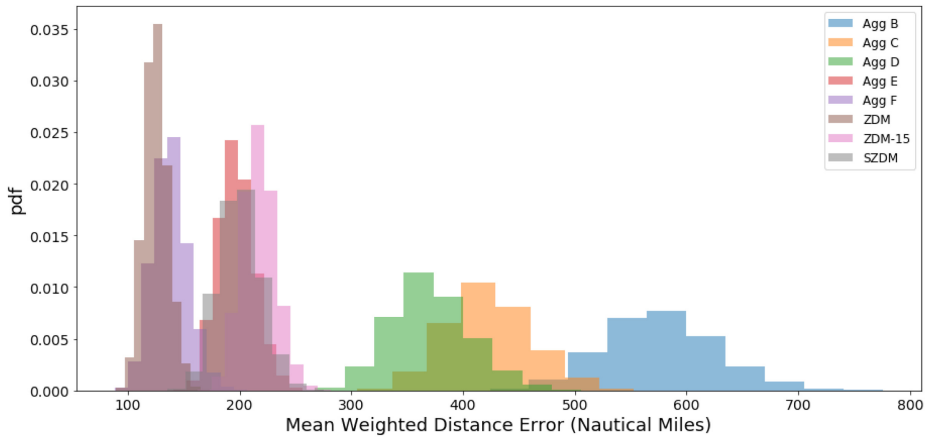Figure 13. Distribution of mean weighted distance errors for all aggregations

Aggregation	No. of zones	Average monthly error (mean zonal demand)	Average monthly error (median zonal demand)
A	1	5.4	5.3
B	2	7.4	7.3
C	8	11.9	11.3
D	15	14.3	12.8
E	44	19.4	15.5
F	212	25.5	20.9
ZDM	36	23.5	19.1

Interestingly, implementing the ZDM corresponds to a larger than expected volume-based aggregation error. With only 36 zones ZDM has a volume-based error higher than Aggregation E which has 44 zones. This suggests deliberate clustering based on geographic proximity does not correspond to improvements in deterministic demand volume modeling. Combined with the results in Section 1 it is clear there is no "one size fits all" aggregation which provides both a low distance-based error and a low volume-based error. Aggregation methods that perform well in one error metric perform correspondingly worse in the other.

4.3 Stochastic volume-based aggregation error
A comparison of stochastic demand models was conducted using probability distributions fit to each zone in Aggregations A–F and SZDM. Afshartous *et al.* (2009) noted that USCG SAR operations could be represented as Poisson processes, and Hornberger *et al.* (2021) continued this approach by constructing a stochastic model using Poisson and Gamma-Poisson distributions. To provide an even comparison, stochastic demand models were constructed for all the quadrat aggregation methods and SZDM based upon the training data set. For each zone, a Gamma-Poisson distribution was used to represent the monthly SAR demand except for instances when the associated parameters made such a fit infeasible; the model defaulted to representing workload with a Poisson distribution when a Gamma-Poisson was deemed inappropriate.

The analysis sought to identify the effectiveness of each stochastic demand model by comparing their abilities to accurately model the demand levels observed in the test data set. Demand models would be considered to be more effective at modeling demand if most of the observed demand levels within the test set fell within anticipated volumes, thereby having fewer months of SAR demand considered *outliers*. To this end, the analysis adopted the Tukey method (Tukey, 1977) for outlier detection, whereby Tukey fences are established at 1.5 times the interquartile range below the 25th percentile and above the 75th percentile of each zone's distribution. Monte Carlo simulations of 1,000 months were constructed for each zone, and the Tukey fences were computed based upon the simulation results.

Table 4 shows the results of this comparison. The general trend observed in these results is consistent with the finding from the literature; predictions of demand volume tend to lose accuracy as the number of zones increase. As the region is decomposed into more zones, the number of represented cases within each zone decreases. This phenomenon leads to more zones whose distributions have smaller within-zone variances. As a result, SAR cases occurring in new locations or at a moderately heightened volume can be registered as outlier events in demand models where the number of zones is higher but not in demand models where the variability of the region is aggregated into only a few zones.

Regarding the stochastic ZDM's effectiveness compared to quadrat models, Table 4 indicates the intentional aggregation of events based on locational clusters offers greater stochastic demand modeling accuracy than an agnostic grid decomposition. The stochastic ZDM's performance adhered to the overall trend, having more identified outliers than Aggregation C and fewer than Aggregation E, but it was able to outperform Aggregation D despite containing the same number of zones. These findings suggest that, in relation to stochastic demand modeling, the method of aggregation combined with the degree of aggregation influences the accuracy of the resulting model.

5. Conclusion

The method used to aggregate spatiotemporal demands affects the outcome of location models built using the aggregated data, thus an understanding of the impacts of aggregation methods is fundamental. While numerous location optimization models require, for tractability, aggregation of spatiotemporal data the creation of aggregation zones is also vital for projection of future demands as the basis of stochastic location analyses. We have presented a framework comparing both static and stochastic spatiotemporal aggregation models, using both a distance-based aggregation error metric, an event magnitude weighted distance-based aggregation error metric and a volume-based aggregation error metric. We further applied this framework to test six quadrat aggregation models of varying fidelities, and two zonal based models, using historical SAR data from a massive scale region possessing highly variable demands. As expected aggregations with greater fidelity tend to

Aggregation	No. of zones	No. of outlier months
A	1	3
B	2	5
C	8	17
D	15	68
E	44	251
F	212	368
SZDM	15	48

Table 4.
Number of outlier months, as defined by demands exceeding Tukey fences, over 1,000 simulated months

reduce the distance-based aggregation error but increase volume-based aggregation error. These results effectively point out that there is no single silver bullet aggregation methodology, rather the researcher must understand the trade-offs in volume- versus distance-based accuracy between methods.

Comparing the quadrat and zonal aggregation methodologies, we find that while the zonal methods observed similar trends in aggregation effects they tended to diminish the impact of those effects. Both ZDM and SZDM approaches exhibited considerably less distance-based error compared to quadrat models with similar numbers of zones. This indicates that location error owing to aggregation can be reduced by constructing aggregation zones based on the inherent clustering of data points through methods such as k-means clustering. Although the ZDM was outperformed by nearly all the quadrat models relative to deterministic demand error, the SZDM displayed less volume-based error compared to quadrat models with similar number of zones. Notably, throughout the duration of this analysis, SZDM consistently outperformed Aggregation D despite both containing 15 zones. This insight offers stochastic zonal aggregation methods such as SZDM as an enticing alternative when aggregation is desired.

Based on the results of our aggregation analysis, we propose the following as potential exploratory efforts. Zonal techniques based on hierarchies and clustering techniques seem very promising, additional research on the impacts of clustering techniques could be fruitful. Additionally combining these zonal techniques, with their associated reduced location errors, with a lower fidelity aggregation model to project region level demands may be useful. Finally, a study examining possible nonlinear dynamic effects on the resulting output of location models as a result of changes in aggregation method may be informative.

References

Afshartous, D., Guan, Y. and Mehrotra, A. (2009), "Us Coast guard air station location with respect to distress calls: a spatial statistics and optimization based methodology", *European Journal of Operational Research*, Vol. 196 No. 3, pp. 1086-1096.

Ai, Y.-F., Lu, J. and Zhang, L.-L. (2015), "The optimization model for the location of maritime emergency supplies reserve bases and the configuration of salvage vessels", *Transportation Research Part E: Logistics and Transportation Review*, Vol. 83, pp. 170-188.

Akbari, A., Eiselt, H.A. and Pelot, R. (2018a), "A maritime search and rescue location analysis considering multiple criteria, with simulated demand", *INFOR: Information Systems and Operational Research*, Vol. 56 No. 1, pp. 92-114.

Akbari, A., Pelot, R. and Eiselt, H.A. (2018b), "A modular capacitated multi-objective model for locating Maritime search and rescue vessels", *Annals of Operations Research*, Vol. 267 Nos 1/2, pp. 3-28.

Anselin, L., Cohen, J., Cook, D., Gorr, W. and Tita, G. (2000), "Spatial analyses of crime", *Criminal Justice*, Vol. 4 No. 2, pp. 213-262.

Araz, C., Selim, H. and Ozkarahan, I. (2007), "A fuzzy multi-objective covering-based vehicle location model for emergency services", *Computers and Operations Research*, Vol. 34 No. 3, pp. 705-726.

Armstrong, R.D. and Cook, W.D. (1979), "Goal programming models for assigning search and rescue aircraft to bases", *Journal of the Operational Research Society*, Vol. 30 No. 6, pp. 555-561.

Azofra, M., Pérez-Labajos, C., Blanco, B. and Achutegui, J. (2007), "Optimum placement of sea rescue resources", *Safety Science*, Vol. 45 No. 9, pp. 941-951.

Curtis, A. and MacPherson, A. (1996), "The zone definition problem in survey research: an empirical example from New York state", *The Professional Geographer*, Vol. 48 No. 3, pp. 310-323.

Dark, S.J. and Bram, D. (2007), "The modifiable areal unit problem (MAUP) in physical geography", *Progress in Physical Geography: Earth and Environment*, Vol. 31 No. 5, pp. 471-479.

Eck, J., Chainey, S., Cameron, J. and Wilson, R. (2005), *Mapping Crime: Understanding Hotspots*, National Institute of Justice.

Erdemir, E.T., Batta, R., Spielman, S., Rogerson, P.A., Blatt, A. and Flanigan, M. (2008), "Optimization of aeromedical base locations in New Mexico using a model that considers crash nodes and paths", *Accident Analysis and Prevention*, Vol. 40 No. 3, pp. 1105-1114.

Fotheringham, A.S., Densham, P.J. and Curtis, A. (1995), "The zone definition problem in location-allocation modeling", *Geographical Analysis*, Vol. 27 No. 1, pp. 60-77.

Francis, R.L. and Lowe, T.J. (1992), "On worst-case aggregation analysis for network location problems", *Annals of Operations Research*, Vol. 40 No. 1, pp. 229-246.

Francis, R.L., Lowe, T.J., Rayco, M.B. and Tamir, A. (2009), "Aggregation error for location models: survey and analysis", *Annals of Operations Research*, Vol. 167 No. 1, pp. 171-208.

Francis, R.L., Lowe, T.J., Tamir, A. and Emir-Farinas, H. (2004), "Aggregation decomposition and aggregation guidelines for a class of minimax and covering location models", *Geographical Analysis*, Vol. 36 No. 4, pp. 332-349.

Hornberger, Z.T., Cox, B.A. and Lunday, B.J. (2021), "Optimal heterogeneous search and rescue asset location modeling for expected spatiotemporal demands using historic event data", *Journal of the Operational Research Society*, pp. 1-19.

Jelinski, D.E. and Wu, J. (1996), "The modifiable areal unit problem and implications for landscape ecology", *Landscape Ecology*, Vol. 11 No. 3, pp. 129-140.

Karatas, M., Razi, N. and Gunal, M.M. (2017), "An ILP and simulation model to optimize search and rescue helicopter operations", *Journal of the Operational Research Society*, Vol. 68 No. 11, pp. 1335-1351.

Kodinariya, T.M. and Makwana, P.R. (2013), "Review on determining number of cluster in k-means clustering", *International Journal*, Vol. 1 No. 6, pp. 90-95.

Li, B. and Szeto, W. (2019), "Taxi service area design: formulation and analysis", *Transportation Research Part E: Logistics and Transportation Review*, Vol. 125, pp. 308-333.

Openshaw, S. (1984), *The Modifiable Areal Unit Problem*, Geo Books, Norwich.

Rajendran, S. and Zack, J. (2019), "Insights on strategic air taxi network infrastructure locations using an iterative constrained clustering approach", *Transportation Research Part E: Logistics and Transportation Review*, Vol. 128, pp. 470-505.

Razi, N. and Karatas, M. (2016), "A multi-objective model for locating search and rescue boats", *European Journal of Operational Research*, Vol. 254 No. 1, pp. 279-293.

Tukey, J.W. (1977), *Exploratory Data Analysis*, Vol. 2, Reading, MA.

US Coast Guard (2013), The US coast guard addendum to the United States National Search and Rescue supplement to the International Aeronautical And Maritime Search and Rescue Manual. COMDTINST M16130.2F.

US Coast Guard (2014), Fourteenth coast guard district search and rescue plan. CGD14INST M16130.1A.

Corresponding author

Bruce Cox can be contacted at: bruceacox1@gmail.com