

UC Office of the President

Recent Work

Title

The p-Compact-regions Problem

Permalink

<https://escholarship.org/uc/item/20t5m3kd>

Journal

Geographical Analysis, 46(3)

Authors

Li, Wenwen
Goodchild, Michael F
Church, Richard

Publication Date

2014-07-14

Peer reviewed

The p-Compact-regions Problem

Wenwen Li¹, Richard L. Church², Michael F. Goodchild²

¹GeoDa Center for Geospatial Analysis and Computation, School of Geographical Sciences and Urban Planning, Arizona State University, Tempe, AZ, USA, ²Department of Geography, University of California, Santa Barbara, CA, USA

The p-compact-regions problem involves the search for an aggregation of n atomic spatial units into p-compact, contiguous regions. This article reports our efforts in designing a heuristic framework—MERGE (memory-based randomized greedy and edge reassignment)—to solve this problem through phases of dealing, randomized greedy, and edge reassignment. This MERGE heuristic is able to memorize (ME of MERGE) the potential best moves toward an optimal solution at each phase of the procedure such that the search efficiency can be greatly improved. A dealing phase grows seeded regions into a viable size. A randomized greedy (RG of MERGE) approach completes the regions' growth and generates a feasible set of p-regions. The edge-reassigning local search (E of MERGE) fine-tunes the results toward better objectives. In addition, a normalized moment of inertia (NMI) is introduced as the method of choice in computing the compactness of each region. We discuss in detail how MERGE works and how this new compactness measure can be seamlessly integrated into different phases of the proposed regionalization procedure. The performance of MERGE is evaluated through the use of both a small and a large p-compact-regions problem motivated by modeling the regional economy of Southern California. We expect this work to contribute to the regionalization theory and practice literature. Theoretically, we formulate a new model for the family of p-compact-regions problems. The novel NMI introduced in the model provides an accurate, robust, and efficient measure of compactness, which is a key objective for p-compact-regions problems. Practically, we developed the MERGE heuristic, proven to be effective and efficient in solving this nonlinear optimization problem to near optimality.

Introduction

The problem of regionalization involves generating larger and fewer regions from smaller but more numerous spatial units to serve a specific purpose. For example, a normative region might be a subdivision of an area with a population size large enough to effectively perform functions such as the administration of elections. An analytical region is formed by grouping zones together

Correspondence: Wenwen Li, School of Geographical Sciences and Urban Planning, Arizona State University, COOR Hall 5644, 975 S. Myrtle Ave., Tempe, AZ 85287-5302
e-mail: wenwen@asu.edu

Submitted: July 7, 2012. Revised version accepted: April 17, 2011.

to facilitate visualization on maps or the interpretation of information at a different scale (Wise, Haining, and Ma 2001). A regionalization problem, within different contexts, can also be termed zonation or districting, relating it to the modifiable areal unit problem (Openshaw 1977). As such, it has been studied and applied in many domains. For example, emergency medical service providers aim to partition a service area into regions to provide satisfactory coverage (Marianov and ReVelle 1996; Alsalloum and Rand 2006; Sorensen and Church 2010). In police deployment, efficient and compact patrol areas are desired (Curtin, Hayslett-McCall, and Qiu 2010). In conservation planning, decision makers may want to identify regions that maximize species coverage and the preservation of selected species (Church, Stoms, and Davis 1996), perhaps combined with the minimization of cost (Moilanen 2007) and clustering (Fischer and Church 2003). Partitioning an area into compact electoral districts is one possible approach to prevent political gerrymandering (Morrill 1976; Young 1988; Pang et al. 2010). In human geography, a single set of reporting zones for projects such as the census is no longer sufficient to meet researchers' needs (Young, Martin, and Skinner 2009), resulting in flexible aggregation becoming increasingly important.

In these applications, compactness of a region is often listed as a desirable goal. First, compactness implies maximum accessibility to all parts of a region. Second, a compact region is likely to be homogeneous, sharing common attributes and properties. Angel, Parent, and Civco (2010, p. 441) acknowledged compactness as "one of the most intriguing and least-understood properties of geographic shapes." Metrics of compactness often compare a shape to a circle of the same area as the shape on the grounds that a circle minimizes the sum of squared distances of all parts of its area from its centroid (Wentz 2000). However, many regionalization studies do not include an explicit compactness measure in their optimization models, relying instead on indirect objectives, such as minimizing travel distances within a region (Wie and Chai 2004) and minimizing the perimeter lengths of regions (Fischer and Church 2003) for simplicity. Such approaches may also allow a mixed integer programming (MIP) model to be formulated so that optimal solution for the problem is easier to obtain.

Literature

The family of regionalization problems involves the aggregation of basic spatial units into groups to meet a range of criteria, including maximizing both intragroup homogeneity and intergroup heterogeneity. Openshaw (1977) defines a general model of the zone-design process by aggregating n zones (i.e., spatial units) into p contiguous regions, where p is usually a small fraction of n . A closely related variant is the location-allocation problem (Goodchild 1979, 1988), which seeks to locate p facilities optimally and simultaneously to allocate distributed demand to them, forming p zones. Models of this type have been applied in many fields and are used to solve a variety of problems, ranging from public facility location-allocation (Hess et al. 1965; Church 1990; Church and Scaparra 2007; Church and Murray 2009) to business service center allocation (Zoltners and Sinha 1983; Aboolian, Sun, and Koehler 2009).

Regional compactness has been considered a very important characteristic in such regionalization problems as political redistricting and sales territory assignment (Adams and Bischof 1994; Duque, Ramos, and Suriñach 2007). Weaver and Hess (1963) presented a methodological framework for political districts that maximizes district compactness while keeping the population totals in each district equal. The compactness measure they used is an indirect measure of the moment of inertia (MI), defined as the sum of all basic units in a region of the products of

population density times the square of the distance from the basic unit to the centroid of the region. This measure of compactness provides only a crude approximation to the true MI, which is more accurate in measuring the shape of an areal object. Another of its weaknesses is that it cannot guarantee spatial contiguity. In 1976, Morrill proposed the use of a computer-based capacity-constrained location-allocation model to generate contiguous congressional districts such that the aggregated travel in each district to its center is minimized (Morrill 1976). This procedure is relatively simple, being based on straight-line travel rather than on a real transportation network, and not considering similarity among the atomic geographic units in the aggregation process. Similar works in the literature include those of Mehrotra, Johnson, and Nemhauser (1998), Bozkaya, Erkut, and Laporte (2003), Ricca and Simeone (2008), and Altman and McDonald (2010).

To provide a basis for formulating the p-compact-regions problem, we review the p-regions problem. This problem deals with the aggregation of n spatial units into p contiguous groups while minimizing the sum of intraregional dissimilarity (Duque, Church, and Middleton 2011). (Note that the “unit” we mention throughout the article means the areal unit.) Several possible ways exist for formulating the p-regions problem as an integer programming problem. Duque, Church, and Middleton (2011) presented three different model formulations for this problem and compare the efficacy of each model. One of the forms is called Order^{RM}, based on the rationale that units are assigned to a given region in a specific order to prevent cycles and, at the same time, to ensure contiguity. Order^{RM} is formulated with the following parameters:

- I = the set of units
 - i, j = indices used to refer to specific units, where $i, j \in I$
 - k = index used to refer to regions, where $k = 1, 2, 3, \dots, p$
 - o, O = index and set of contiguity order
 - $N_i = \{j \mid \text{area } j \text{ is adjacent to area } i\}$
 - S_{ij} = a measure of dissimilarity between unit i and unit j
- Utilizing the following decision variables:

$$t_{ij} = \begin{cases} 1, & \text{if units } i \text{ and } j \text{ belong to the same region } k \text{ where } i < j \\ 0, & \text{otherwise} \end{cases}$$

$$x_i^{ko} = \begin{cases} 1, & \text{if unit } i \text{ is assigned to region } k \text{ in order } o \\ 0, & \text{otherwise} \end{cases}$$

the ORDER model can be defined as

$$\text{Minimize: } \sum_i \sum_{j>i} S_{ij} t_{ij}, \tag{1}$$

subject to

$$\sum_i x_i^{k0} = 1, \tag{2}$$

$$\sum_{k=1}^p \sum_o x_i^{ko} = 1, \tag{3}$$

$$x_i^{ko} \leq \sum_{j \in N_i} x_j^{k(o-1)}, \forall i \in I, \quad \forall k = 1, 2, 3, \dots, p, \quad \text{and } \forall o \in O, \tag{4}$$

$$t_{ij} \geq \sum_o x_i^{ko} + \sum_o x_j^{ko} - 1, \forall k = 1, 2, 3, \dots, p, \quad \forall i \in I, \quad \text{and } \forall j \in I, \text{ where } i < j, \tag{5}$$

$$x_i^{ko} \in \{0, 1\}, \forall i \in I, \forall k = 1, 2, 3, \dots, p \text{ and } o \in O, \tag{6}$$

$$t_{ij} \in \{0, 1\}, \forall i \in I, \forall j \in I, \quad \text{where } i < j. \tag{7}$$

This model minimizes the dissimilarity between each pair of areas assigned to the same region. Constraint (2) ensures that each region is assigned exactly one unit as the seed with order zero. Constraint (3) ensures that each unit is assigned to exactly one region under one order of assignment. Constraint (4) ensures that a unit cannot be assigned to region k as order o unless a neighboring area j has been assigned to region k as order $o - 1$. Thus, this constraint ensures that any area assigned to region k is a neighbor of the root area of region k or is a neighbor of an area that has been assigned to region k with a lower order of closeness. Altogether, constraint (4) ensures that each region is a contiguous set of units. These constraints were inspired by the work of Cova and Church (2000). Duque, Church, and Middleton (2011) tested two other models (Tree and Flow), reporting that all model forms are of limited use in solving a p-regions problem to optimality beyond a problem involving 50 areas. This is an issue that has been raised in a previous work as well (Bixby 2002). The remaining issue here is that only contiguity is considered, and any concern for compactness is missing. Because of this, we propose a variant of the p-regions problem—the p-compact-regions problem—that involves the solution for a compact set of regions. The p-compact-regions problem involves maximizing compactness of resultant regions rather than minimizing dissimilarity between each pair of areas within the same region.

Problem statement and model formalization

In this section, we will introduce the formulation and common constraints of p-compact-regions problem, and the integration of normalized moment of inertia (NMI) as the novel compactness measure.

The p-compact-regions problem: a general model

The p-compact-regions problem involves a single objective that attempts to optimize a measure of overall compactness of the p-regions. In essence, we can assume a formulation of the p-regions problem, like that given previously where we substitute the dissimilarity objective for one that involves maximizing compactness. To solve this problem, a robust and computationally feasible compactness measure is the key to defining such a model. Li, Goodchild, and Church (2013) reviewed both direct and indirect compactness measures and proposed a new way to compute compactness of a shape based on the NMI; they proved that this approach provides a robust and computationally efficient compactness measure. Therefore, we adopt their approach here by defining a region Z_k 's compactness index $C(Z_k)$ as

$$C(Z_k) = \frac{A_{Z_k}^2}{2\pi I_{Z_k}^G}, \tag{8}$$

where Z_k is the set of units that are assigned to region k , A_{Z_k} is the area of Z_k , and $I_{Z_k}^G$ is the second MI of a region about an axis perpendicular to it and passing through its centroid G . $C(Z_k)$ ranges from zero, in the case of an infinitely extended shape, to one, in the case of the most compact figure, a circle.

Given this definition of compactness of a region Z_k , we define the sum of the NMI of each of the p -regions as the compactness objective of the p -compact-regions problem:

$$\text{Maximize: } \sum_{k=1}^p C(Z_k). \tag{9}$$

Given that the range of $C(Z_k)$ is between (0, 1], the range of the compactness objective value is (0, p]. This can be thought of as a different objective to the p -regions model with constraints (2)–(7), which focuses on compactness, where contiguity is a required property for each region and all areas' units must be assigned to exactly one region. To derive $C(Z_k)$ in the objective function, it is essential to derive the MI of region Z_k , which is $I_{Z_k}^G$.

Compactness—MI $I_{Z_k}^G$

The measure of compactness $I_{Z_k}^G$ in equation (8) could be based on pure geometry; as such, it is known as the area-weighted MI. $I_{Z_k}^G$ could also represent the MI weighted by the distribution of a nongeographical property, such as population. Li, Goodchild, and Church (2013) proposed a new trapezium-based approach to compute the area-weighted MI. In this article, we introduce an extension of their MI definition to cope with the needs of different applications. We start by defining $I_i^{G_i}$, the MI of a basic unit, which is the building block to computing the MI of a region Z_k . Mathematically,

$$I_i^{G_i} = \frac{A_i \int d^2 \rho_i(da_i) da_i}{\int \rho_i(da_i) da_i}, \tag{10}$$

where A_i is the area of a basic unit i , da is an infinitesimal part of area A_i , $\rho_i(da_i)$ is the density function of some property of i , such as population, G_i is unit i 's density-weighted centroid, and d is the distance from da_i to G_i . Assuming that a property is evenly distributed over each basic unit, namely $\rho(da_i)$ is a constant, ρ_i , $I_i^{G_i}$ can be simplified to

$$I_i^{G_i} = \rho_i \int d^2 da_i, \tag{11}$$

where $I_i^{G_i}$ is equal to ρ times the area-weighted MI of the unit. When $\rho_i = 1$, $I_i^{G_i}$ can be computed based on pure geometry (see Li, Goodchild, and Church 2013 for details of the calculation).

Because the computation of MI is additive (Li, Goodchild, and Church 2013), the MI of a region Z_k can be derived by summing the MI of all its contained units plus the MI introduced by this aggregation. Fig. 1 shows a partial region Z_k composed of three basic units, the boundaries of which are drawn in dashed lines. This region has a neighboring unit i , with area A_i , density ρ_i , and centroid G_i . A neighborhood is determined based on the sharing of common edges (i.e., rook contiguity). If ρ_{Z_k} is the mean density of region Z_k and G is region Z_k 's density-weighted centroid, after aggregating unit i into Z_k , the new MI of Z_k becomes




Figure 1. Location change of region Z_k 's centroid for computing the density-weighted MI.

$$I_{Z_k}^{G'} = I_{Z_k}^G + I_i^{G_i} + \rho_{Z_k} A_{Z_k} d_1^2 + \rho_i A_i d_2^2, \tag{12}$$

$$(x', y') = \left(\frac{\rho_{Z_k} A_{Z_k} x + \rho_i A_i x_i}{\rho_{Z_k} A_{Z_k} + \rho_i A_i}, \frac{\rho_{Z_k} A_{Z_k} y + \rho_i A_i y_i}{\rho_{Z_k} A_{Z_k} + \rho_i A_i} \right), \tag{13}$$

$$\rho'_{Z_k} = \frac{\rho_{Z_k} A_{Z_k} + \rho_i A_i}{A_{Z_k} + A_i}, \tag{14}$$

where $G'(x', y')$ is the new density-weighted centroid after a unit i is added to region Z_k . Its coordinates can be computed with equation (13). d_1 is the Euclidean distance between G' and G_i , and d_2 is the Euclidean distance between G_i and G' . Region Z_k 's new density after adding unit i , ρ'_{Z_k} , can be computed using equation (14). Substituting $I_{Z_k}^{G'}$ in equation (12) for $C(Z_k)$ in equation (9) provides a formal definition of the compactness objective, where compactness can be computed as a region grows or shrinks when member units are reassigned. The computation, although additive, is based on knowing the centroid of the region through a nonlinear process, according to equations (12)–(14). Therefore, the p-compact-regions problem, as formulated in equations (2)–(7) and (9), is a nonlinear integer programming problem.

Specialization of a p-compact-regions problem

According to different application needs, the p-compact-regions problem may incorporate additional constraints or objectives. The following is a list of a few possible constraints required in addressing real-world problems:

- (1) Size constraint: a size constraint could require the resultant regions to be of equal size or require any region to be bigger or smaller than a certain size.
- (2) Physiographic constraint: boundaries of regions cannot cross major physiographic features, such as mountain range barriers. This constraint is often used in ecological/economic regionalization.

- (3) Hierarchical constraint: in multilevel geographic systems, frequently the hierarchy of political boundaries must be preserved.
- (4) A constraint on intra- and interzonal flow could require that intrazonal flow not exceed a certain fraction of the interzonal flow, or vice versa.
- (5) Deviation constraint for a desired target: this requires that the difference between the value of a property for a region and an ideal value be limited. The sum of such a deviation also could be minimized.

MERGE: A heuristic framework to solve the p-compact-regions problem

Given the nonlinear nature of the p-compact-regions problem and our desire to solve large-sized problems, a heuristic framework termed MERGE has been developed to solve the p-compact-regions problem. MERGE has a sequence of stages: initial growth, growth completion, and local search. It starts with p seeds as beginning locations about which p -regions are constructed (Brookes 1997; Church et al. 2003). Although a good heuristic algorithm should not be affected by the selection of the seeds, an appropriate seed selection strategy tailored to a specific application is conducive to the solving of regionalization problems. One strategy is to select several units at random as seeds; this strategy has been used in many applications (Vickrey 1961; Thoreson and Liittschwager 1967; Openshaw 1977; Rossiter and Johnston 1981). The random-seed selection strategy is suitable for small problems without many constraints, such as the hexagon application we present later. As for other more complex problems, constraints inherent in a problem may provide useful guidance for selecting seeds. For instance, the transportation analysis zone (TAZ) aggregation problem presented in the “Application” section has a constraint that no local commercial centers (each one is represented as a set of basic units) can be assigned to different regions, and no region may contain more than one local commercial center. In this case, the local commercial centers naturally become the seeds for growing regions. Once guidelines for selecting seeds are identified for a specific problem, one could either manually or randomly select seeds.

After seed selection, MERGE starts the initial growth stage through a “dealing” procedure that ensures that each region “grows” to a viable size. Stage II of MERGE applies a randomized greedy algorithm to the partial regions resulting from Stage I to create a set of regions that meets all of the constraints (2)–(7). Thus, the completion of Stage II is a feasible solution consisting of a set of p contiguous regions. Stage III of the heuristic allows for possible backtracking in order to find improvements by using a form of simulated annealing (SA) (Kirkpatrick, Gelatt, and Vecchi 1983) applied to possible edge reassignment. “Edge reassignment” here means the reassignment of basic units on the edge of regions. A basic SA algorithm starts with an initial solution and then randomly selects a feasible neighboring solution. Here, a neighboring solution is a solution generated by reassigning a basic unit on the edge of a region to one of its neighboring regions (neighbor detection is based on rook contiguity) without violating the contiguity constraint. If this neighboring solution is an improvement over the current solution, it is used to replace the existing one, and the search continues. Otherwise, the neighboring solution is accepted using an acceptance probability function $P = e^{-\Delta T}$, also known as the Boltzmann equation (Arkeryd 1972), where

Δ is the variation of the evaluation criterion and T is a global time-varying variable representing temperature. T is initialized at T_0 and is decreased after a set number of iterations by a cooling rate α . When T reaches some predefined value ϵ , the iterative process stops, and the current best solution is accepted.

Fig. 2 shows the flowchart of MERGE. To ensure that all regions grow to a viable size, we initially deal a fixed number M of basic units to each region, assigning them on the basis of the greatest improvement in the objective function. We call this “dealing” because it involves something akin to a card dealer who deals, in a clockwise fashion, one card at a time to each player until each player has a set number of cards. Here, we deal one of the best units to a seeded region and continue dealing until each seeded region has one of the best neighboring units assigned to it. We then deal a second set of units and then a third set and a fourth set and so on until each seed cluster has a size of $M + 1$ units (M assigned units plus the original seed unit). Without this dealing step, some regions may fail to grow beyond two or three units if adding another basic unit decreases rather than increases the objective function value. This dealing step is deceptively simple, but extraordinarily powerful, as a region could be quite awkward in size and shape when it contains only two or three basic units, and adding a unit at this point often makes a cluster less compact rather than more compact. But, allowing all regions to grow (by adding basic units) beyond an awkward shape, such as an enclave, helps to tune the heuristic toward generating a high percentage of compact regions.

After the dealing step, the randomized greedy algorithm is applied globally over the p-regions. For each region, the N best plans for adding a basic unit are selected by looking at all possible units that can be added. Then, one of the N is randomly selected. At this point, a basic unit may be eligible to be added to several different regions. After the plans for all of the regions are identified, the assignment to a region that results in the largest objective function value improvement is selected. Each region has a selected candidate, and that candidate is one of the N best possibilities. This means that a degree of randomness is associated with the choice of which basic unit is added to a region and which region is selected for the next assigned basic unit. This is a form of randomized adaptive search. After adding the selected basic unit i , the values and plans for the other regions are updated as necessary. To support efficient execution of this algorithm, MERGE introduces a memory module to keep track of a region’s current status, its best N growth plans, and the potential status of a region after executing the best N growth plans for all regions. The status of a region is represented by a number of variables, including the identifications codes of all basic units within each region, the region’s compactness index, its centroid, area, MI, and other properties. Before the assignment of any unit is complete, the compactness index of each region must be updated in memory if the availability of a basic unit is changed. Instead of updating the status of all regions after a basic unit is assigned, only the regions that are affected need to be updated. For example, when unit i is assigned at an iteration, the status of the region to which it is assigned must be updated by recalculating the compactness index, the quantitative value for each constraint, and the objective function value. In addition, the algorithm needs to identify new candidate units that can improve the composite objective function of the region. Other regions that may be affected are those that already specify unit i in their candidate lists. For such regions, if region i becomes unavailable, another candidate unit must be identified to replace it in the region’s candidate list (if any are adjacent and unassigned).

To improve the regionalization process further, an edge-reassigning approach was designed to tune the results at Stage III (as Fig. 2b shows). Its input is the set of region partition plans obtained from Stage II. The strategy is to select randomly a unit on the edge of a region and then to consider allocating it to an adjacent region that has the largest ΔObj ; that is, the greatest overall objective function improvement or the least objective function value decrease, as expressed in the following equation:

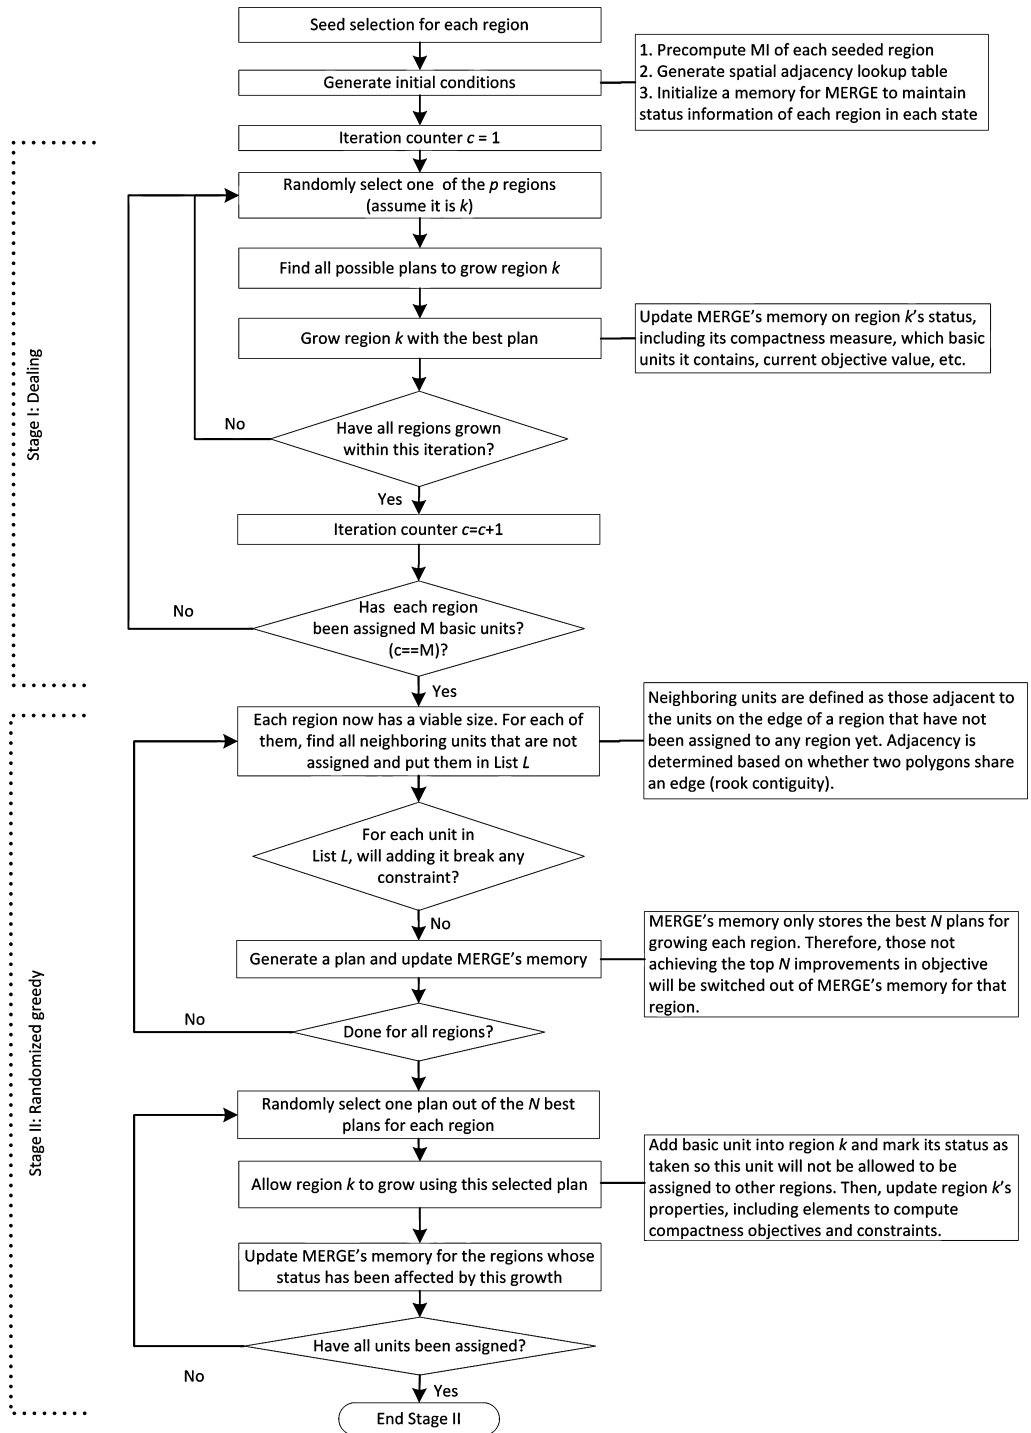


Figure 2. (a) Flowchart of Stage I (dealing) and Stage II (randomized greedy) in MERGE. (b) Flowchart of Stage III (local search) in MERGE.

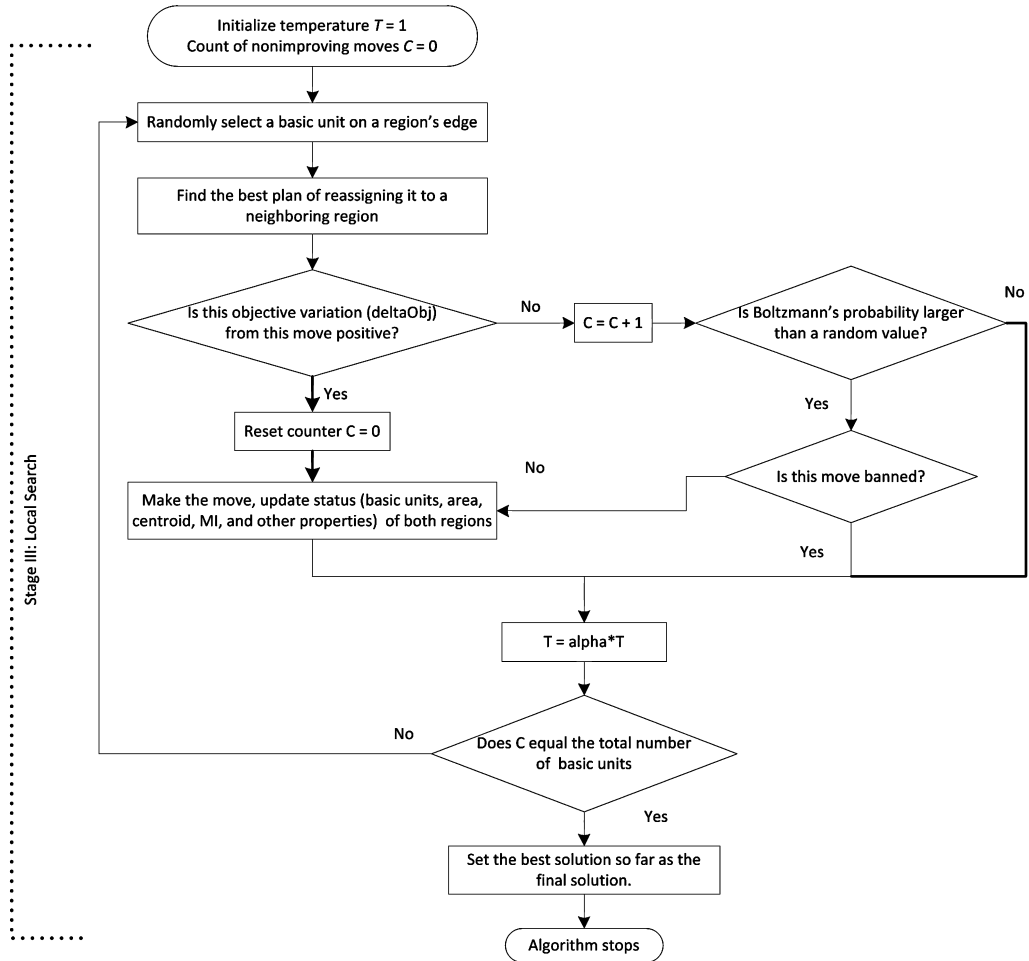


Figure 2. Continued

$$\Delta Obj = Obj_{new}(P_k) + Obj_{new}(P_u) - Obj(P_k) - Obj(P_u). \tag{15}$$

The selected move is made when ΔObj is positive. When ΔObj is negative or zero, it is accepted with a probability given by the Boltzmann equation. In our approach, we make three improvements to what would be a simple SA algorithm: First, instead of selecting a feasible neighboring solution randomly, the procedure selects a locally best reassigning for the unit that improves the speed of approach to an optimal solution. As the local search continues (T decreases), the solution tends to converge to a stable state, which one hopes is near-optimal. According to the Boltzmann equation, the possibility of allowing a nonimproving move decreases as T decreases. Because of this mechanism, the nonimproving moves mostly occur at an early stage of the local search when the probability of accepting a nonimproving move is greatest when T is large.

Second, as downhill nonimproving moves are allowed, the process may repeatedly visit the same solution, especially at a later stage of the edge-reassigning process when the heuristic is nearing completion. Because almost all of the possible improving moves have been tried at this point and because the nonimproving moves can be accepted only with a very low probability,

when a nonimproving move suddenly is allowed, the very next move would most likely be the reverse move that produces an increase in the objective value. Then, the process bounces among a small set of solutions. To avoid this pitfall, a list of banned moves can be introduced, like that found in a tabu search (Glover 1989, 1990). When an uphill move is made, its reverse move is added to the banned set in the tabu list. Even if its reverse downhill move is chosen using Boltzmann's probability, this reverse move is not allowed when the move is banned. When the list is full, the oldest move is dropped from the list (often called tabu tenure) and becomes unbanned.

Third, the stopping condition is strengthened. In a general SA method, the algorithm stops when T reaches a predefined value. However, this process could be very time-consuming when the final value of T is set to be very small. To make efficient use of computational effort, MERGE is stopped when all of the last K (i.e., the total number of basic units) sequentially selected candidate moves are nonimproving, whether or not they are executed. When this condition is encountered, the best determined solution is selected as the final result.

In MERGE, the assignment or reassignment of a basic unit cannot break the contiguity constraint. During the region-growing phases (Stage I and Stage II), a basic unit that is added to a region is selected from a set of all the unassigned basic units adjacent to the region. Therefore, contiguity for a region is preserved without the need for a specific strategy. In the local search phase, however, removing one basic unit from a region potentially may violate the contiguity constraint by leaving a residual region split into two unconnected parts. To address this issue, one could check whether a path exists between any of the unit pairs in a region, but this would be a very time-consuming process. The strategy we adopt here is to select one basic unit within a region and to find its adjacent basic units within the same region until no more units can be added. If the size of the identified unit set is less than the size of the region, then one or more basic units are unreachable after the edge reassignment. Therefore, the edge unit cannot be reassigned.

Finally, the design of this heuristic is such that it is to be repeated hundreds, if not thousands, of times. After each of its executions, the best solution found is saved for further analysis. This process represents a design feature of semi-greedy and greedy randomized adaptive search procedure (GRASP)-based processes where the idea is to introduce random choices between several of the best performing options as the heuristic progresses. This design feature allows a wide variety of very good starting feasible solutions to be generated, under which the improved Stage III is used to test whether local transitions lead to better local optima or even to a global optima.

A simple example is given to demonstrate how MERGE works. Fig. 3 depicts two partial regions (P_1 and P_2) and three unassigned units (1–3). The entire region is an 8×4 rectangle; unit 2 is a 1×4 rectangle; unit 3 is a 3×2 rectangle (as is region P_2); and unit 1 is a right isosceles triangle (horizontal and vertical edges length). The objective of this p-compact-regions problem is simplified to address only compactness. Suppose regions P_1 and P_2 have grown from seed units through the dealing phase to their current state. Fig. 3 lists their MI, area, and objective function value (obtained from equation [8]). d_1 is the distance from the centroid of a region in its current state to the centroid of this region once the unit is added. d_2 is the distance from the centroid of a unit to the centroid of a region after the unit is added. For example, if unit 3 is added to region P_2 , the centroid of region P_2 falls on the midpoint of the border segment shared between region P_2 and unit 3; therefore, $d_1 = d_2 = 1$, as listed in Table 1. Because region P_1 and unit 3 as well as region P_2 and unit 1 are not neighbors, no moving plans will be made for these (region, unit) pairs.




Figure 3. An example of region growth via the MERGE algorithm.

Table 1. Distance of Potential Moves

Distance	Region P_1	Plans	Region P_2	Plans
Unit 1	d_1 0.2694	Plan B: Assign unit 1 to region P_1	N/A	N/A
	d_2 1.8856			
Unit 2	d_1 0.5142	Plan C: Assign unit 2 to region P_1	d_1 0.8944	Plan D: Assign unit 2 to region P_2
	d_2 1.8100		d_2 1.3416	
Unit 3	N/A	N/A	d_1 1.0000	Plan A: Assign unit 3 to region P_2
			d_2 1.0000	

Two plans exist for growing region P_2 : adding unit 3 to P_2 (Plan A) or adding unit 2 to P_2 (Plan D). According to equation (8), the MI of new region P'_2 after executing Plan A becomes

$$I_{P'_2} = I_{P_2} + I_3 + A_{P_2}d_1^2 + A_3d_2^2 = 6.5 + 6.5 + 6 \times 1^2 + 6 \times 1^2 = 25.$$

Correspondingly, the NMI for region P'_2 yields

$$Obj_{new(P'_2)} = \frac{A_{P'_2}^2}{2\pi I_{P'_2}} = \frac{(6+6)^2}{2 \times 3.14159 \times 25} = 0.9167.$$

Thus, the objective function value increase brought by Plan A is 0.0352 [$\Delta(Obj) = 0.9167 - 0.8815 = 0.0352$]. It is the overall best plan in comparison to all other plans, including Plan B for adding unit 1 to region P_1 [$\Delta(Obj) = 0.0279$], Plan C for adding unit 2 to region P_1 [$\Delta(Obj) = -0.007$], and Plan D for adding unit 2 to region P_2 [$\Delta(Obj) = -0.2229$]. If Plan B is executed, P_1 's compactness objective term becomes the highest (0.9549) over all plans; but, because we count the greatest increase in the objective function value instead of the highest objective function value that can be achieved, Plan A still is considered the best plan so far. Using the proposed randomized greedy algorithm, one plan is randomly selected from the best N growth plans for each region, and then the best plan over all selected plans is executed. Suppose Plans A and D are ranked as the top two plans for growing region P_2 , Plans B and C are ranked as the top two plans for growing region P_1 , and Plans D and C are randomly selected to grow these two regions, respectively. Because Plan C

increases the objective function value more than Plan D does, at this round, unit 2 will be added to region P_2 . Although Plan A is the overall best plan, it is not picked due to the randomization process. Instead, another relative best move is made. In this way, we allow not only regions to grow toward more optimal solutions but also downhill moves to help escape from local optima.

The preceding example explains how randomized greedy (Stage II of MERGE) works; we now demonstrate how the proposed local search (Stage III of MERGE) works. Suppose after Stage II that all basic units are assigned. Region P_1 contains its original part and units 1 and 2; region P_2 contains its original part and unit 3. During the first round, each basic unit is scanned, and only unit 2 is identified as a unit on the edge. It has one possible move: remove from region P_1 and merge into region P_2 . Suppose that, after this move, P_1 becomes (a 4×4 square) and P_2 becomes (a 44×4 square), where the objective function value can be calculated as

$$Obj(P_1) + Obj(P_2) = 0.9316 + 0.9167 = 1.8483 < Obj(P'_1) + Obj(P'_2) = 0.9549 + 0.9549 = 1.9098.$$

Thus, this move leads to a better solution; therefore, it is allowed. After this round, each of the basic units is inspected again, and because no moves will produce an increase in the overall objective function value, the algorithm stops. In this case, this solution (two regions with a 4×4 square each) is an optimal solution with regard to the goal of generating two compact regions.

Applications

The previous section introduces the MERGE algorithm in detail and how it can be used to solve the p -compact-regions problem. In this section, we introduce two compactness-driven applications of the p -compact-regions problem with different constraints. We begin with a small problem having a known optimal solution to prove the ability of MERGE to achieve optimality. We then introduce a large practical p -compact-regions problem, the solution of which is an essential building block to modeling economic processes. We also compare MERGE with other popular heuristics, including a common greedy algorithm, tabu search (Glover 1989), and GRASP (Feo and Resende 1995), to demonstrate the superiority of MERGE. The algorithms were implemented in Python and tested on a 2×2.93 GHz Quad-Core Intel Xeon CPU, 16G DDR3 memory, and a Mac OS X Lion 10.7.4.

An equal-sized p -compact-regions problem

Fig. 4 shows a hypothetical study area composed of 168 equal-sized units, each of which is an equilateral triangle. These triangles are connected at their vertices, and neighboring triangles share an edge and two vertices. The objective of this p -compact-regions problem is to generate seven ($p = 7$) equal-sized compact regions. That is to say, each region should contain 24 ($168/7$) basic units. The most compact region composed from 24 such triangles is a regular hexagon, whose compactness index is 0.99 using the MI-based measure. Given that hexagonal tiling is the most efficient way to partition a two-dimensional plane into equal areas (Goodall 1987; Hales, Sarnak, and Pugh 2000), the optimized solution of this p -compact-regions problem is the seven adjacent regular hexagons (in different gray shades) illustrated in Fig. 4.

This problem is relatively small, with less than 200 basic units. However, its optimized solution is not easy to obtain due to the many possible growing plans that bring the same improvement to its objective function value. This case is especially true when many partial regions are symmetric. For example, when a partial region contains two units (e.g., units 150 and 151), the plans of adding 149 or 132 or 152 or 163 yield the same improvement to the objective




Figure 4. Triangular lattice for generating $p = 7$ equal-sized compact regions and an illustration of an optimized solution.

function value. Therefore, they all are the best theoretically, and any plan could potentially be selected by the randomized procedure. However, if the plan of aggregating unit 152 is chosen, the regionalization process results in a nonoptimal solution. The proposed local search algorithm in MERGE shows its advantage to adjust a nonoptimal solution to an optimal one in this case.

MERGE starts by randomly selecting seven basic units as seeded regions. Such seeds are used to initialize the process. Once they have served the purpose of being sites to initially coalesce surrounding units, however, these seeds are allowed to be moved to a neighboring region during the edge-reassignment process (Stage III). The local search takes the plans generated through the randomized greedy algorithm as input. The MERGE heuristic was executed 1,000 times for this problem, and Fig. 5 is a plot of the number of times solutions are generated that exhibit a specific average compactness value among the seven regions that fall within a specific




Figure 4. *Continued*

average compactness interval. The intervals of 0.01 in size range from 0.79 to 0.99. The solid gray and black bars represent the distribution of average compactness values by executing Stage II with and without a dealing stage, respectively. The bar filled with a left-slanting line represents the distribution of average compactness values after the end of Stage III. With the dealing heuristic, the randomized greedy algorithm generated feasible plans with an objective function value ranging from 0.89 to 0.99, which improved the results obtained by executing Stage II only. Among these improved feasible solutions, about 5% (value = 0.99) are optimal solutions, and about 15% of the feasible plans of Stage II had objective values of 0.98 (near-optimal). After applying the proposed local search algorithm (Stage III), all of the final plans become optimal; that is, each plan resulted in seven regular hexagons, as shown in Fig. 4. This experiment demonstrates the ability of MERGE to achieve global optimal solutions for a somewhat simple p-compact-regions problem for which a known optimal solution exists.

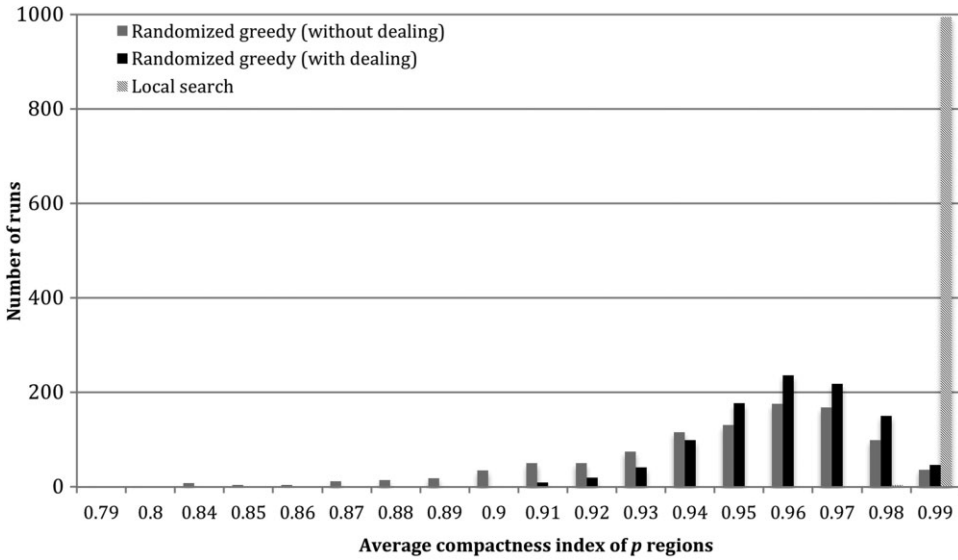


Figure 5. Distribution of the mean compactness index of the seven regions over 1,000 runs.

A large practical p-compact-regions problem

In this section, we demonstrate the solution using MERGE on a real-world p-compact-regions problem with multiple constraints. Our goal is to aggregate the 4,109 TAZs in six Southern California counties (Los Angeles, Riverside, San Bernardino, Orange, Imperial, and Ventura; see Fig. 6) to approximately 100 regions, the maximum number considered computationally feasible for a large-scale spatial econometric model involving land use change and transportation investment (Anas and Liu 2007). We choose to maximize the overall compactness of generated regions and, at the same time, to satisfy the following constraints: (1) each region must be spatially coterminous; (2) the boundaries of regions must not cross major physiographic features or county boundaries; (3) the intrazonal traffic generated within a region must be smaller than a certain percentage of all traffic generated by that region; (4) a local commercial center made up of one or more TAZs cannot be split into two regions; and (5) each region cannot be assigned more than one commercial center. A commercial center is defined as a set of spatially contiguous TAZs, for which the employment density and total employment of all TAZs must be greater than some threshold. This is a nonlinear problem in nature, and its size and model type make it difficult to solve optimally for all but the smallest of sizes.

We applied MERGE to the 4,109 TAZs with a set of 100 manually selected seeds. This seed set comprised 49 local commercial centers and 51 TAZs containing junctions of major roads. The same set of seeds was used for all computer runs. We set the number of TAZs dealt in the first phase of MERGE to $K = 10$, thus accounting for roughly a quarter of all TAZs assigned. We suggest that some experimentation with K may be needed in other applications of this heuristic if the values of n and p are significantly different from those in this application. In the randomized greedy phase, N (the number of plans from which to select randomly when growing a region) was set to three. MERGE was run 1,000 times: Stage I (dealing) and Stage II (randomized greedy) were run together to first generate a set of 1,000 feasible p-region solutions; then Stage III (local search) started with these 1,000 feasible solutions in a search for improved solutions.




Figure 6. An overview of the 4,109 TAZs in the study area.

We also compared the performance of MERGE to two popular heuristics, GRASP and tabu search, to demonstrate its superiority. A GRASP algorithm used in a local search phase always makes the reassignment that brings the greatest positive improvement in the objective function value and continues until no further improvement is possible. Nonimproving moves are not allowed in the improvement phase of GRASP. A tabu search, in comparison, has three strategies. First, it uses a tabu list to prevent cycles by forbidding reverse moves for a number of iterations. This list is dynamic, and eventually an old move can be considered again as a candidate move. Second, tabu allows up to k nonimproving moves to overcome local optima. Third, a tabu search always seeks to make the move that is the best greedy move (the best improvement or the least damaging). A tabu search stops when all possible moves are either infeasible (no more moves are allowed, and the allowed number of nonimproving moves has been reached) or all feasible moves are taboo. Both GRASP and tabu were applied to the initial feasible solutions resulting from Stage II. As 1,000 feasible solutions were generated by Stage I of MERGE, the GRASP and tabu heuristics were run 1,000 times each.

Fig. 7 shows the distribution of objective function values over 1,000 runs for dealing and randomized greedy at Stages I and II. The objective function value equals the sum of the compactness index over p -regions (in this case, $p = 100$). Hence, an objective function value

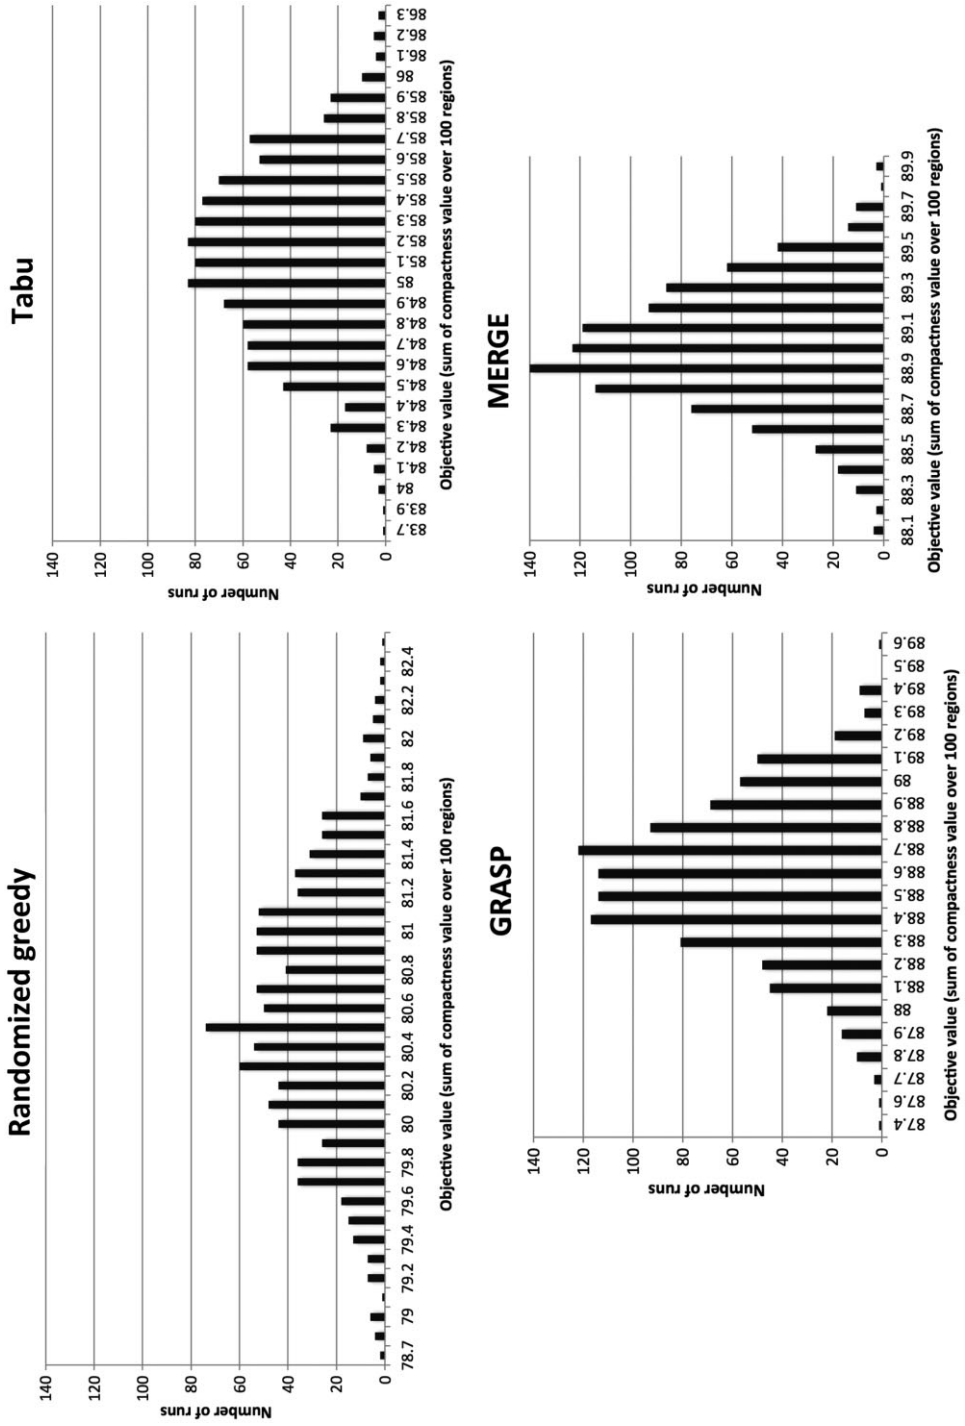


Figure 7. Distribution of the objective function values of the p-compact regions generated by tabu, GRASP, and MERGE.

equaling 88 signifies that the mean compactness index of the p-regions is 0.88 (1/100 of the objective value), which has one as its upper bound and zero as its lower bound. The objective function values range between 78.7 and 82.4 after the completion of the randomized greedy phase in MERGE, as the first cluster of columns shows in Fig. 7. This dispersion is caused by the random choices introduced in MERGE to determine which regions grow and in what order. Out of the 1,000 runs, 578 (57.8%) solutions have an objective function value that is better than what would be generated by a nonrandomized greedy algorithm (i.e., 80.5). Clearly, multiple runs with the randomized greedy procedure in MERGE are capable of producing better starting solutions.

Based on these solutions, tabu, GRASP, and MERGE were, respectively, applied to further improve the region partition plans. The second cluster of columns, with objectives ranging from 83.7 to 86.3, are those obtained by tabu. Tabu-85 in the figure indicates that the maximal length of the tabu list is 85, the same parameter value used in the max-p-regions problem (Duque, Anselin, and Rey 2012), and in that used by Ricca and Simeone (2008). Although tabu achieves an average of 5.64% objective function improvement over the randomized greedy approach, it is still far worse than MERGE or GRASP. The reasons why tabu fails to make much of an improvement in this regionalization problem are multifaceted. First, occurrence of a cycling move in a seed-growing approach is unlikely because distant basic units are unlikely to interact with each other. Hence, when moves are selected to generate more compact regions at each step, they tend to generate very different moves instead of the reverse moves of the ones just made. Second, although tabu allows nonimproving moves, the number is a predefined value; tabu does not have the ability to adjust the number dynamically according to the size and complexity of the problem. This explains why the tabu strategy does not work well for the given problem with a large solution space.

In comparison with tabu (with objective function values ranging from 83.7 to 86.3), GRASP and MERGE perform much better. GRASP generates solutions with values ranging between 87.4 and 89.4. GRASP benefits from the randomized greedy approach, which generates a wide variety of solutions by not always accepting the best moves. Therefore, although GRASP does not allow nonimproving moves during the local search phase, it is still able to achieve an average of 9.9% objective function value increase vis-à-vis the initial solutions. Over all algorithms, the best solutions were found with MERGE. It results in solutions with values ranging from 88.1 to 89.9, and on average, it produced an improvement of over 10.5% in objective function value over that of the initial solutions. Although some overlaps exist in objective function values with GRASP, for an individual run using the same initial plan, MERGE outperformed GRASP in 960 out of 1,000 cases. In addition, only MERGE found the best solution ever identified (objective value = 89.9). This superior performance of MERGE during the local search phase can be attributed to the following strategies: (1) the probabilistic acceptance of nonimproving moves, which is the essence of the SA method; (2) the adoption of only the best move for a selected basic unit, which keeps the algorithm from going too far from the optimal solution and helps the algorithm to converge quicker; (3) the stop condition that intelligently ends the search procedure when an optimal or near-optimal plan is found; and (4) the list of banned nonimproving moves to prevent the algorithm from getting stuck in a local optima. Fig. 8 shows the best region partition plan (objective = 89) found by MERGE. Visually, most regions are round or nearly round. Other features include the following: (1) the shape of some regions is adversely affected by boundary effects, such as the elongated region south of the Santa Monica Mountains in Los Angeles County and (2) some regions are inside other solution regions. For example, five small regions are surrounded by one big region in Riverside County. This is because most population and economic activity in Riverside County occurs within the five small regions. From west to east, they are the

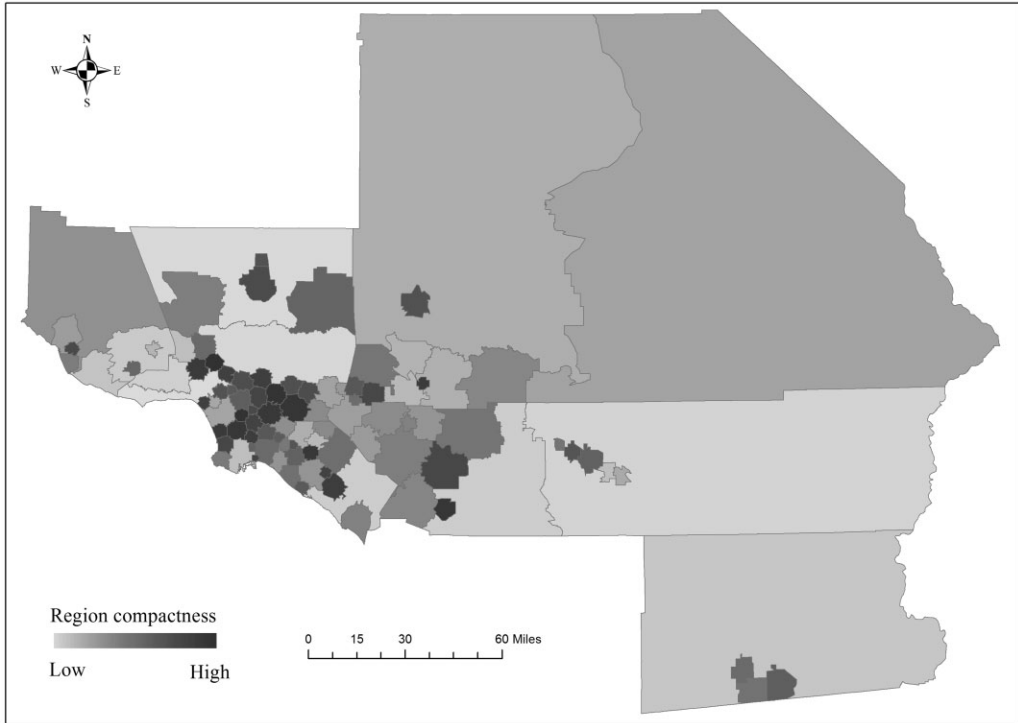


Figure 8. The best region partition plan obtained by MERGE.

cities of Palm Springs, Cathedral City, Palm Desert, La Quinta, and Indio. The rest of the region in the county is mountainous or desert or has low population density (with fewer economic activities and associated traffic). Instead, most of these activities occur between the suburban areas and the city centers. Therefore, as a whole, the rest of the area is partitioned into one region.

In terms of time efficiency, on average, 76 seconds was required to complete a run with tabu search, 154 seconds for GRASP, and 200 seconds for MERGE. Although a little bit slower, MERGE is able to find better solutions than those found by tabu and GRASP.

Discussion and conclusion

In this article, we introduce a new formulation of the p-compact-regions problem, aiming to generate a fixed number of compact regions from a larger number of basic units. This formulation integrates the NMI as its compactness measure because this measure’s effectiveness and efficiency have been proven (Li, Goodchild, and Church 2013). This index is able to measure compactness based not only on pure geometry but also on weighting by some property, such as population. Although a number of researchers (Dantzig and Ramser 1959; Burns et al. 1985; Laporte 2009) in the field of management science have proposed measuring compactness as the total length of the network or travel routes within a defined region, this measure is potentially suitable for dividing a network into regions rather than aggregating areal units into regions. Our work, in comparison, provides an accurate measure of compactness based on a region’s shape (and possibly weighted by some attribute like population density) and therefore has the potential to be applied to a much broader application domain.

The other major contribution of this work is the design of the MERGE heuristic that is able to solve efficiently large practical regionalization problems. The following new strategies contribute to its superior performance: a dealing heuristic is particularly helpful in preventing a region from getting stuck at a state that is very small but compact (this is the foundation for generating a good solution); a dynamically updatable memory module activated throughout the randomized search procedure avoids repeated unproductive computations, such that the algorithm can be very efficient, even for a large problem; a SA-inspired local search allows MERGE to exploit a larger search space to obtain a set of final solutions from which to select the best solution.

The results reported here are counter to those of Ricca and Simeone, who state that “simulated annealing is not able to reach a good local optimum for compactness” (Ricca and Simeone 2008, p. 1425). Potentially, three reasons exist for this discrepancy. First, Ricca and Simeone used the common SA approach without attempting to fine-tune it for their political districting problem. Second, their cooling rate ($\alpha = -0.998$) is faster than that set in MERGE for this article ($\alpha = -0.999$). Because of this, Ricca and Simeone’s algorithm tends to make fewer downhill moves, taking less iteration to reach the final temperature or stopping condition. This is likely to result in poor local optima as the potential of SA to sample the surrounding neighborhood of solutions is restricted. In terms of the compactness measure, Ricca and Simeone used a distance model to compute the compactness index, which is only a rough estimation of a region’s compactness. This is quite different from a MI metric, which is a function of the square of distances to the exact centroid. Unfortunately, they do not describe how starting feasible solutions are generated. An arbitrary region partition plan as a starting solution also may affect the performance of the local search procedure. However, to understand this or to draw any real conclusion about it requires further investigation.

Finally, the proposed model and solution framework can be extended to solve problems with two or more objectives. A sensible extension of the p-compact-regions problem is to a multiobjective compact-p-regions problem, which includes both the compactness objective and the similarity measure introduced in Duque, Church, and Middleton (2011). The new composite objective function can be defined as the weighted combination of a region’s compactness and the dissimilarity within a region:

$$\text{Maximize: } w_1 \times C(Z_k) - w_2 \sum_i \sum_{j>i} s_{ij} t_{ij}, \tag{16}$$

where $i, j \in Z_k$ and $w_1, w_2 \in [0, 1]$. Because the second term involves minimizing dissimilarity, it is preceded by a negative sign, which translates it into an equivalent maximizing objective. Solving this problem with the MERGE framework makes sense for future work. Application of the p-compact-regions problem to one of the political districting problems, where the lack of compactness is often seen as a surrogate measure of the presence of gerrymandering, also is worth exploring.

Acknowledgements

This research was funded in part by the University of California Multi-campus Research Initiative (MRPI). Dr. Sungbin Cho of the Southern California Association of Governments, who provided GIS data for our urban regionalization application, is also acknowledged.

References

- Aboolian, R., Y. Sun, and G. J. Koehler. (2009). "A Location-Allocation Problem for a Web Services Provider in a Competitive Market." *European Journal of Operational Research* 194(1), 64–77.
- Adams, R., and L. Bischof. (1994). "Seeded Region Growing." *IEEE Transactions of Pattern Analysis and Machine Intelligence* 16, 641–47.
- Alsalloum, O. L., and G. K. Rand. (2006). "Extensions to Emergency Vehicle Location Models." *Computers and Operations Research* 33, 2725–43.
- Altman, M., and M. McDonald. (2010). "The Promise and Perils of Computers in Redistricting." *Duke Journal of Constitutional Law and Public Policy* 5, 69–111.
- Anas, A., and Y. Liu. (2007). "A Regional Economy, Land Use, and Transportation Model (RELU-TRAN): Formulation, Algorithm Design and Testing." *Journal of Regional Science* 47(3), 415–55.
- Angel, S., J. Parent, and D. Civco. (2010). "Ten Compactness Properties of Circles: Measuring Shape in Geography." *Canadian Geographer* 54(4), 441–61.
- Arkeryd, L. (1972). "On the Boltzmann Equation." *Archive for Rational Mechanics and Analysis* 45(1), 1–16.
- Bixby, R. E. (2002). "Solving Real-World Linear Programs: A Decade and More of Progress." *Operations Research* 50, 3–15.
- Bozkaya, B., E. Erkut, and G. Laporte. (2003). "A Tabu Search Heuristic and Adaptive Memory Procedure for Political Districting." *European Journal of Operational Research* 144(1), 12–26.
- Brookes, C. J. (1997). "A Parameterized Region Growing Programme for Site Allocation on Raster Suitability Maps." *International Journal of Geographical Information Science* 11, 375–96.
- Burns, L. D., R. W. Hall, D. E. Blumenfeld, and C. F. Daganzo. (1985). "Distribution Strategies that Minimize Transportation and Inventory Costs." *Operations Research* 33(3), 469–90.
- Church, R. L. (1990). "The Regionally Constrained P -Median Problem." *Geographical Analysis* 22(1), 22–32.
- Church, R. L., and A. T. Murray. (2009). *Business Site Selection, Location Analysis, and GIS*. Hoboken, NJ: Wiley.
- Church, R. L., and M. Scaparra. (2007). "Protecting Critical Assets: The R -Interaction Median Problem with Fortification." *Geographical Analysis* 39(2), 129–46.
- Church, R. L., D. M. Stoms, and F. W. Davis. (1996). "Reserve Selection as a Maximal Covering Location Problem." *Biological Conservation* 76(2), 105–12.
- Church, R. L., R. A. Gerrard, M. Gilpin, and P. Stine. (2003). "Constructing Cell-Based Habitat Patches Useful in Conservation Planning." *Annals of the Association of American Geographers* 93(4), 814–27.
- Cova, T. J., and R. L. Church. (2000). "Contiguity Constraints for Single-Region Site Search Problems." *Geographical Analysis* 32(4), 306–29.
- Curtin, K. M., K. Hayslett-McCall, and F. Qiu. (2010). "Determining Optimal Police Patrol Areas with Maximal Covering and Backup Covering Location Models." *Networks and Spatial Economics* 10, 125–45.
- Dantzig, G. B., and J. H. Ramser. (1959). "The Truck Dispatching Problem." *Management Science* 6, 80–91.
- Duque, J. C., R. Ramos, and J. Suriñach. (2007). "Supervised Regionalization Methods: A Survey." *International Regional Science Review* 30, 195–220.
- Duque, J. C., L. Anselin, and S. J. Rey. (2012). "The Max- P -Regions Problem." *Journal of Regional Science* 52, 397–419.
- Duque, W. C., R. L. Church, and R. S. Middleton. (2011). "The P -Region Problem." *Geographical Analysis* 43, 104–26.
- Feo, T. A., and M. G. C. Resende. (1995). "Greedy Randomized Adaptive Search Procedures." *Journal of Global Optimization* 6, 109–33.
- Fischer, D., and R. L. Church. (2003). "Clustering and Compactness in Reserve Site Selection: An Extension of the Biodiversity Management Area Selection Model." *Forest Science* 49(4), 555–65.
- Glover, F. (1989). "Tabu Search—Part I." *ORSA Journal on Computing* 1, 190–206.

Geographical Analysis

- Glover, F. (1990). "Tabu Search—Part II." *ORSA Journal on Computing* 2, 4–32.
- Goodall, B. (1987). *The Penguin Dictionary of Human Geography*. London: Penguin.
- Goodchild, M. F. (1979). "The Aggregation Problem in Location Allocation." *Geographic Analysis* 11(3), 240–55.
- Goodchild, M. F. (1988). "Location-Allocation on a Microcomputer: The PLACE Package." In *Desktop Planning: Microcomputer Applications for Infrastructure and Services Planning and Management*, 166–71, edited by P. W. Newton, M. A. P. Taylor and R. Sharpe. North Melbourne: Hargreen.
- Hales, T., P. Sarnak, and M. Pugh. (2000). "Advances in Random Matrix Theory, Zeta Functions, and Sphere Packing." *Proceedings of the National Academy of Sciences of the United States of America* 97, 12963–64.
- Hess, S. W., J. Weaver, H. Siegfeldt, J. Whelan, and P. Zitlau. (1965). "Nonpartisan Political Redistricting by Computer." *Operations Research* 13, 998–1006.
- Kirkpatrick, S., C. Gelatt, and M. Vecchi. (1983). "Optimization by Simulated Annealing." *Science* 220, 671–80.
- Laporte, G. (2009). "Fifty Years of Vehicle Routing." *Transportation Science* 43, 408–16.
- Li, W., M. F. Goodchild, and R. L. Church. (2013). "An Efficient Measure of Compactness for Two-Dimensional Shapes and Its Application in Regionalization Problems." *International Journal of Geographical Information Science* 27(6), 1227–50.
- Marianov, V., and C. ReVelle. (1996). "The Queuing Maximal Availability Location Problem: A Model for the Siting of Emergency Vehicles." *European Journal of Operational Research* 93, 110–20.
- Mehrotra, A., E. L. Johnson, and G. L. Nemhauser. (1998). "An Optimization-Based Heuristic for Political Districting." *Management Science* 44(8), 1100–114.
- Moilanen, A. (2007). "Landscape Zonation, Benefit Functions, and Target-Based Planning: Unifying Reserve Selection Strategies." *Biological Conservation* 134(4), 571–79.
- Morrill, R. L. (1976). "Redistricting Revisited." *Annals of Associations of American Geographers* 66, 548–66.
- Openshaw, S. (1977). "A Geographical Solution to Scale, and Aggregation Problems in Region-Building, Partitioning, and Spatial Modeling." *Transactions of the Institute of British Geographers, New Series* 2, 169–84.
- Pang, S., H. He, Y. Li, T. Zhou, and K. Xing. (2010). "An Approach of Redistricting Based on Simple and Compactness." *Advances in Swarm Intelligence* 6145, 415–24.
- Ricca, F., and B. Simeone. (2008). "Local Search Algorithms for Political Districting." *European Journal of Operational Research* 189(3), 1409–26.
- Rossiter, D., and R. J. Johnston. (1981). "Program GROUP: The Identification of All Possible Solutions to a Constituency-Delimitation Problem." *Environment and Planning* 13, 231–38.
- Sorensen, P., and R. L. Church. (2010). "Integrating Expected Coverage, and Local Reliability for Emergency Medical Services Location Problems." *Socio-Economic Planning Science* 44(1), 8–18.
- Thoreson, J. D., and J. M. Liittschwager. (1967). "Computers in Behavioral Science: Legislative Districting by Computer Simulation." *Behavioral Science* 12, 237–47.
- Vickrey, W. (1961). "On the Prevention of Gerrymandering." *Political Science Quarterly* 76, 105–10.
- Weaver, J. B., and S. W. Hess. (1963). "Procedure for Nonpartisan Districting: Development of Computer Techniques." *Yale Law Journal* 73, 288–308.
- Wentz, E. A. (2000). "A Shape Definition for Geographic Applications Based on Edge, Elongation, and Perforation." *Geographical Analysis* 32, 95–112.
- Wie, B., and W. Chai. (2004). "An Intelligent GIS-Based Spatial Zoning System with Multiobjective Hybrid Metaheuristic Method." In *Proceedings of the 17th International Conference on Innovations in Applied Artificial Intelligence*, 769–78, edited by B. Orchard, C. Yang and A. Moonis. Berlin and Heidelberg: Springer-Verlag.
- Wise, S., R. Haining, and J. Ma. (2001). "Providing Spatial Statistical Data Analysis Functionality for the GIS User: The SAGE Project." *International Journal of Geographical Information Science* 15, 239–54.

- Young, C., D. Martin, and C. Skinner. (2009). "Geographically Intelligent Disclosure Control for Flexible Aggregation of Census Data." *International Journal of Geographical Information Science* 23(4), 457–82.
- Young, H. P. (1988). "Measuring the Compactness of Legislative Districts." *Legislative Studies Quarterly* 13(1), 105–15.
- Zoltners, A. A., and P. Sinha. (1983). "Sales Territory Alignment—A Review and Model." *Management Science* 29(11), 1237–56.