

Chapter 15

Heuristics for Location Models

Jack Brimberg and John M. Hodgson

15.1 Introduction

Location-allocation problems, due to their mathematical complexity, resist exact solutions for problems of more than moderate size. For this and other reasons, heuristic (approximative) approaches are widely used in solving them. This chapter considers the two seminal streams of heuristic solution procedures, both of which remain in use today often in a somewhat altered form. We begin by outlining the location-allocation problem that originally attracted the development of these approaches.

Consider the following basic scenario: a set of n points is given, either in the continuous plane or serving as nodes in a network, with each expressing a demand for some service. These demands can be met through travel to or from facilities also located at points. The objective is to determine a specified number of facility location points that provide the best possible service to these demand points. The studies considered here explicitly define this problem as determining that set of facility locations that minimizes a sum of demand-weighted distances between the demand points (customers) and their nearest facility. This problem is known as the multi-median or generalized median problem. Commonly, the number of facility locations sought is denoted by p , leading to the alternate and more popular title of p -median problem. Some researchers make a distinction between the continuous space *multisource Weber problem* (Brimberg et al. 2000) and the discrete or network space *p-median problem*; here, we do not.

Classic location problems including the location-allocation type occur in two-dimensional space, which may in turn be depicted in three ways: continuous, discrete, and network space. In the continuous case, referred to as site-generating models

J. Brimberg (✉)

Department of Mathematics and Computer Science, Royal Military College of Canada,
Kingston, ON K7K 7B4, Canada
e-mail: jack.brimberg@rmc.ca

J. M. Hodgson

Department of Earth and Atmospheric Sciences, The University of Alberta,
Calgary, AB T6G 2R3, Canada
e-mail: john.hodgson@ualberta.ca

(see, e.g., Love et al. 1988), locations are given by Euclidean coordinates, distances are calculated endogenously from these coordinates, and facility locations are determined by *solving* for the p best pairs of (x, y) coordinates in the plane. In the discrete case, facility locations are *selected* from a set of potential sites (site selection models) and the distances, or more specifically, the shipping costs between demand points and potential sites are provided exogenously. In the network case, demands are expressed at the vertices of networks, facility locations (in the case of medians) are *selected* from network vertices, and distances are calculated over the shortest paths in the network. The classic papers of focus in this chapter were presented in continuous and network space; in the context of the p -median model, network and discrete space are identical in all practical terms.

Regardless of the space considered, the p -median problem is computationally difficult, having been shown to be NP-hard by Kariv and Hakimi (1979) for networks and Megiddo and Supowit (1984) for the continuous case. Thus, heuristic approaches are required to solve problems of reasonable size. This difficulty occurred to early researchers in location-allocation analysis, and seminal papers considered in this chapter presented the model formulations along with heuristic algorithms to solve them. They presented heuristic algorithms of two types. Cooper's early work with the continuous space problem gave rise to the *alternating locate/allocate* heuristic, which Maranzana adapted to the network space problem. Teitz and Bart dealt with the network problem using a *vertex substitution* heuristic in which vertices are systematically shifted in and out of a trial solution set.

The remainder of the chapter is organized as follows. Section 15.2 of this chapter reviews the pioneering papers of Cooper (1963, 1964) and Maranzana (1964) including their well-known alternating locate/allocate heuristics; we then proceed to the classical vertex substitution heuristic of Teitz and Bart (1968). Section 15.3 examines the impact of these seminal works on later developments. Section 15.4 provides a short discussion of the future direction of research on the p -median problem. Finally, Sect. 15.5 offers some concluding remarks.

15.2 The Classical Contributions

This section will follow the developments of ideas of four seminal papers in the field of location analysis. Each of these papers includes some heuristic methods that have had an impact on the way location problems are solved.

15.2.1 Cooper (1963, 1964)

Cooper (1963) posed the problem of locating a set of *sources* (facilities) in some optimal fashion in order to serve a set of *destinations* (customers) at fixed and known locations. The problem was described in the following general terms: given the

location of each destination, the requirements at each destination, and a set of shipping costs for the region of interest, determine the number of sources, the location of each source, and the capacity of each source.

Alluding to the theoretical difficulties of this problem, Cooper then added two important assumptions: there are no capacity restrictions on the facilities, and unit shipping costs are independent of facility output.

To put the new problem in context, Cooper (1963) reviewed the background literature on single facility problems, touching on the work of Cavalieri, Steiner, and others. Today, we know that this brief review overlooked two papers of major importance, those of Weiszfeld (1937) and Kuhn and Kuenne (1962), which provided a solution method for the single facility problem.

Cooper stated the problem formally, but we have changed the notation for overall consistency. The known customer locations are defined by their Cartesian coordinates

$$(a_i, b_i), \quad i = 1, 2, \dots, n,$$

and the coordinates of the p facilities to be determined are

$$(x_j, y_j), \quad j = 1, 2, \dots, p.$$

Note that although p is now given, we may repeat the analysis for various values of this parameter in order to ultimately determine the ‘best’ number of facilities.

Cooper also indicated that “in addition to not knowing the location of each of the p sources in the minimum cost solution, we also do not know which source is to serve which subset of destinations.” We now term this *allocations*. He introduced a binary variable to deal with these allocations:

$$\alpha_{ij} = 1 \quad \text{if customer } i \text{ is served by facility } j, \quad 0 \text{ if not.}$$

He further introduced a weighting factor w_{ij} relating to the multiplicity of supply trips or service calls, a measure of demand for the service.

Cooper framed the problem mathematically in general terms and then introduced the notion of optimal service being the minimization of a weighted sum of Euclidean distances between the customers and the facilities that serve them. This leads to the following formulation:

$$\text{Min: } \varphi = \sum_{j=1}^p \sum_{i=1}^n \alpha_{ij} w_{ij} [(a_i - x_j)^2 + (b_i - y_j)^2]^{1/2}. \quad (15.1)$$

Setting the first-order partial derivatives of (15.1) to zero with respect to each x_j and y_j provides conditions for a minimum. Thus, after replacing the Euclidean distance terms in (15.1) with

$$D_{ij} = [(a_i - x_j)^2 + (b_i - y_j)^2]^{1/2}$$

and re-arranging, we obtain

$$x_j = \frac{\sum_{i=1}^n \frac{\alpha_{ij} w_{ij} a_i}{D_{ij}}}{\sum_{i=1}^n \frac{\alpha_{ij} w_{ij}}{D_{ij}}} \text{ and } y_j = \frac{\sum_{i=1}^n \frac{\alpha_{ij} w_{ij} b_i}{D_{ij}}}{\sum_{i=1}^n \frac{\alpha_{ij} w_{ij}}{D_{ij}}}, \quad j = 1, 2, \dots, p. \quad (15.2)$$

Each pair of simultaneous equations provides the optimal solution for the single median problem defined for a known subset of customers who are allocated to a particular source. Where p facilities are sought, there will be p allocation groups; hence, p separate single facility problems must be solved for any given set of allocations (or partition of the customer set). Cooper presented an iterative scheme for solving (15.2) with fixed allocations, which is now commonly termed the Weiszfeld procedure, having rediscovered the same iterative scheme first proposed by Weiszfeld (1937) for the single median problem. The procedure simply updates the coordinates of each facility by substituting the values from the latest iteration in the right-hand side of (15.2), and continuing in this fashion until convergence is detected. (The interested reader is referred to Kuhn 1973, and Katz 1974, for convergence studies of the Weiszfeld procedure.) As a starting point for the iterations, Cooper recommended the weighted mean center of the customer points:

$$x_j^0 = \frac{\sum_{i=1}^n \alpha_{ij} w_{ij} a_i}{\sum_{i=1}^n \alpha_{ij} w_{ij}}, \text{ and } y_j^0 = \frac{\sum_{i=1}^n \alpha_{ij} w_{ij} b_i}{\sum_{i=1}^n \alpha_{ij} w_{ij}}, \quad j = 1, 2, \dots, p. \quad (15.3)$$

(Cooper presents these values incorrectly in both the 1963 and 1964 papers, by neglecting the w_{ij} factor in the denominators. His case studies survive this error because all weights are assumed equal to unity.)

Cooper's initial solution approach is straightforward, based on his observation of the crux of the location-allocation relationship: "If, for a set of n destinations and p sources, the location of the sources is known, the determination of the optimal allocations is trivial. It is merely the set of weighted distances [...] that is a minimum. Conversely, if the allocation is fixed, the determination of the optimum location of the sources is merely the exact calculation, with *known* α_{ij} that has been previously described" (Cooper 1964). The problem can thus be solved exactly by examining all possible allocation sets, $\{\alpha_{ij}\}$, and choosing the solution that minimizes (15.1). He presented a test problem, but acknowledged that this approach would not be computationally attractive for what in the 1963 computing environment was quaintly considered to be a large set of customers (>10). Cooper determined that the number of such allocations is the Stirling number of the second kind $S(n, p)$, a number that remains "formidably large" for problems that are considered of modest size today.

Cooper observed that a method for generating a "reasonable" number of facility location sets was required. He suggested considering the n customer sites as potential facility locations, thus reducing the problem to the discrete space form.

The problem is thereby reduced to size ${}_nC_p$, but he acknowledged that the method remained “inadequate for many problems of industrial importance because of the excessive amount of calculation involved”. Moreover, he recognized that limiting facilities to a base set of customer sites might not yield correct allocations.

Having introduced the continuous location-allocation problem and discussed the issues in solving this problem in his 1963 paper, Cooper (1964) turned to the development of heuristics to solve it effectively. He surmised that the problem had no sharp minimum, but rather many alternative or close optima, which “makes feasible the use of heuristic algorithms with a reasonably high probability that a well constructed heuristic will find one of these near optimal solutions.” (Here, we see an interesting distinction between classical and modern viewpoints: today, the existence of multiple local minima is considered to be regrettable, and the goal of heuristics is to come as close as possible to the optimal solution rather than to identify a good one. This goal is facilitated, of course, by the enormous increase in computing power available today.) The paper revisits the definition of the problem, the iterative procedure for the single facility problem (which he terms the “exact” procedure), and the suggestion that a direct solution to the multi-source problem is to minimize ϕ for all possible sets of α_{ij} . In order to treat the problem of very large customer sets ($n \leq 500$) and situations of nonlinear costs (an idea he did not pursue further), Cooper developed several heuristic algorithms. He began by proposing lower bounds for limited cases of the problem, which he used to rank the results of the heuristics, and also determined an obvious upper bound for the problem.

Cooper (1964) presented four basic heuristics that are summarized below. Three of the heuristics assume that the “destination set is a very favored set,” and thus, use subsets of the customer set for locating facilities. The first two of these, upon termination, use the exact (Weiszfeld) procedure to determine the (optimal) continuous space origins with respect to the selected allocation.

A: The Destination Subset Algorithm. This considers all possible subsets of p customers as “sites” at which to locate facilities. This is the method of the 1963 paper and basically involves complete enumeration of all ${}_nC_p$ discrete space solutions to the problem, a reduction from the $S(n, p)$ possible continuous space solutions. Upon termination, a continuous adjustment is applied by using the exact procedure on the best discrete solution. Cooper again states that this procedure does not guarantee an optimal solution (the correct allocation may not be found), and warns that computation time for large problems is prohibitive.

B: The Random Destination Algorithm. Here, p random customers are selected to be facility locations. The algorithm is repeated a number of times and the best solution is retained. Cooper suggests a statistical approach to determine when the algorithm might reasonably be terminated. The procedure provides a continuous adjustment on the final solution as before.

C: The Successive Approximation Algorithm. The destination subset algorithm is run for $p=2$ facilities. The best location for a third facility among the remaining customer sites is determined and then locked into the solution. Additional facilities

are similarly added. This is, after the initial two facilities are located, a greedy constructive algorithm of customer sites. It is unclear why Cooper does not apply the exact algorithm after termination.

D: The Alternate Location and Allocation Algorithm. This elegant heuristic by Cooper continues to be popular to this day. It is based on the simple observation alluded to earlier that the two components of the problem, locate and allocate, are easy to solve in isolation. That is, given the locations of the p facilities, and the fact that there are no capacity restrictions on them, the customers are simply allocated in turn to the facility that provides the lowest cost service. For homogeneous facilities ($w_{ij}=w_j$ for all i, j), this translates to assigning each customer to the closest facility. If the allocations are given, the problem reduces to p independent single facility problems that may be readily solved with the Weiszfeld procedure. The heuristic simply alternates between the two phases until no further improvement is possible. The steps provided by Cooper are summarized in *Algorithm 1*.

Cooper identifies the algorithm as a monotonic-decreasing convergent process that may not converge to the globally optimal solution. In fact, the process only guarantees a local minimum.

Cooper then uses solutions to the destination subset algorithm to demonstrate the lack of a sharp minimum, and reiterates that “it is this relative insensitivity to source location *with correct or near-correct allocations* which makes the use of heuristics feasible in this problem.”

Algorithm 1: Alternating Locate/Allocate (ALT-1)

- Step 1:* Divide the customer set into p subsets of approximately equal size.
- Step 2:* For each subset, apply the exact procedure to determine the optimal facility location.
- Step 3:* Allocate all customers to the closest facility.
- Step 4:* Continue alternating between steps 2 and 3 until there is no allocation change.

Cooper tabulated results for 10 problems of size, $n=30$, $p=3$. For the first time in his experience, the destination subset algorithm arrived at an incorrect allocation. He also used 400 iterations of the random destination algorithm. As expected, the destination subset algorithm generally found the best solutions, but was by far the most computationally demanding of the methods. The random destination approach was next in quality, and much less costly. The successive approximation approach was not satisfactory—it is unfortunate that Cooper did not apply the exact solution procedure to its results for a fairer comparison with approaches *A* and *B*. The alternate location and allocation algorithm (ALT-1) was deemed to be “satisfactory.” Cooper neglected to note that this algorithm actually performed best in three of the trials, even though its statistics were troubled by three spectacular failures.

He further tested the heuristics with 100 problems of size $n=40$, $p=3$, and concluded that: “it is apparent that the best *practical* method of solving large location-

allocation problems is with the use of the random destination algorithm and subsequent improvement by a single calculation of the exact location method.” Again, he does not credit the alternate location and allocation approach, the solution statistics for which were almost identical to the random destination approach.

One questions why Cooper downplayed the abilities of the alternating location and allocation algorithm, which has been passed down as the major contribution of his early work. One further questions why Cooper did not think to combine algorithms B and D to apply the alternating algorithm to the random solutions. Cooper started the alternating heuristic with a rather messy selection of an initial allocation. Scott (1971) modified the algorithm by starting with a random selection of trial facility locations. Later work discovered that the influence of a single poor solution could be overcome by using several such random starts. This has become the common procedure for using the alternating heuristic today. The random multi-start version of the alternating heuristic allows us to obtain several local minima in different regions of the solution space, and thus improves the chances of obtaining a “good” solution.

Cooper’s 1963 and 1964 papers first identified the location-allocation problem in mainstream literature. Moreover, they identify the computational characteristics of the problem, while the second paper provides solution techniques for what is otherwise a very difficult class of problems. The alternate location and allocation algorithm is still often used today. For several years, this popular approach has been dubbed “The Cooper Algorithm” (Scott 1971).

15.2.2 Maranzana (1964)

Maranzana (1964) defined the location-allocation problem on a network space as follows: given, in a network, a set V of n points (referred to as “sinks”) v_1, \dots, v_n , with associated nonnegative weights w_1, \dots, w_n , and a nonnegative, n -dimensional, symmetric distance matrix $[d_{ij}]$, find p sources v_{x_1}, \dots, v_{x_p} among the points in V , and a partition of V into p subsets of sinks V_{x_1}, \dots, V_{x_p} served respectively by the p sources so that

$$\sum_{j=1}^p \sum_{v_i \in V_{x_j}} D_{i,x_j} w_i \quad (15.4)$$

is a minimum, where D_{ij} is the minimum path length from v_i to v_j . (Again we have changed notation for consistency.) The total transport cost is assumed to be proportional to the weighted sum of shortest-path distances given in (15.4).

Maranzana concluded that direct enumeration would be impractical for “the typical problem,” and proposed instead an iterative procedure to solve the problem heuristically. The method alternates between location and allocation phases as in Cooper’s algorithm, except that since we are dealing with a network, the shortest paths between all pairs of nodes must be determined first in a preprocessing step. Maranzana adapted a dynamic programming approach attributed to

Bellman (1958) to accomplish this. The shortest path between a given pair of nodes is determined recursively as the shortest path that uses at most j edges, for $j=1, 2, \dots, n-1$. By using a fixed sequence and constant updates of the shortest paths between all pairs of nodes, Maranzana actually improved the efficiency of Bellman's algorithm.

Maranzana also noted that a separate routine was required to find the "center of gravity" of any subset Q of nodes, which he defined as a vertex v_j of V that provides the minimum weighted sum of shortest path lengths between itself, acting as a facility, and the vertices of Q , acting as its customers. (This should not be confused with the median of set Q , since v_j is not restricted to Q .) To find the center of gravity, Maranzana simply evaluated each vertex and chose the best one. The steps of his heuristic may be outlined as shown in *Algorithm 2*.

Algorithm 2: Alternating Locate/Allocate (ALT-2)

- Step 1:* Select p trial facility sites arbitrarily from the n vertices of V to specify the current location set.
- Step 2:* Partition the n vertices by assigning each customer to its nearest facility in the current location set.
- Step 3:* Determine a "center of gravity" for each subset in the partition.
- Step 4:* If the center of gravity is the same as the current location of the facility for each subset, stop (the current location set with associated partition is the final solution); else update the current locations to the new centers of gravity and return to Step 2.

Maranzana proved that the sequence of solutions generated by the algorithm is monotone non-increasing by showing that the allocation and location phases (Steps 2 and 3) may only improve the current solution. He then provided a simple numerical example to show that the procedure can converge to a non-optimal (local) minimum. Using a second simple example, he demonstrated the difficulty that may arise when the center of gravity (Step 3) is non-unique; that is, different decision rules for breaking ties may lead to very different solutions. To circumvent the above difficulties, Maranzana, unlike Cooper, who treated the alternating heuristic as being suited to a single application, suggested that "with a computer it is feasible to carry out the procedure on a number of initial selections so that one can be assured of arriving at a good solution." Finally, the algorithm was applied to problems of two and three facilities in a case study of 158 Italian cities, each given a hypothetical weight that appears to have been related roughly to city size. It is interesting that computation time was mainly devoted to the calculation of shortest path distances on the network, a problem that would be magnified later on for practical applications on much larger networks. Finally, we note that Maranzana seems to have assumed in his procedure that the optimal facility sites are located at the vertices of the network, a result that would be proven coincidentally by Hakimi (1964).

15.2.3 Teitz and Bart (1968)

Teitz and Bart (1968) addressed the problem of solving what they called the “generalized vertex median of a weighted graph,” which today we call the “uncapacitated p -median problem on a network.” Specifically, they considered “the problem of choice of location of p sources of unconstrained capacity from among n destinations having fixed demands and located at nodes of a network.” They acknowledged that their problem was essentially the same as that investigated by Hakimi (1964) and Maranzana (1964), stating that their concern was with alternative methods of solution.

The problem is thus defined on an edge and vertex weighted graph, G . Each vertex v_i is weighted by a weight w_i and each i - j edge by the shortest path distance D_{ij} . The distance matrix \mathbf{D} of G is an $[n \times n]$ symmetric matrix of shortest path distances between all pairs of vertices v_i, v_j . The weighted distance matrix \mathbf{R} , asymmetric for differentially weighted vertices, is defined as

$$[r_{ij}] = [w_i D_{ij}]. \quad (15.5)$$

The single vertex median solves

$$r_k = \min \{r_1, r_2, \dots, r_n\}, \quad (15.6)$$

where

$$r_j = \sum_{i=1}^n r_{ij}, \quad j = 1, \dots, n. \quad (15.7)$$

The generalized vertex median problem can be developed as follows: let V_p be some subset containing exactly p vertices of G . In an n -vertex graph there will be $\binom{n}{p}$ possible such subsets, indexed V_p^m ; $m = 1, 2, \dots, \binom{n}{p}$. For each such subset, we may construct a submatrix \mathbf{R}_p^m of \mathbf{R} by adjoining all columns of \mathbf{R} for which the corresponding column vertices are contained in V_p^m . If facilities are limited to vertices in V_p^m , each customer v_i will be served by that facility in V_p^m for which r_{ik} is a minimum. The total weighted distance r_m for the V_p^m set of facilities is

$$r_m = \sum_{i=1}^n r_{ik}, \quad (15.8)$$

where in each row of \mathbf{R}_p^m , $k(=k(i))$ is the facility for which r_{ij} is minimized. The general vertex p -median of G is defined as some V_p^{m*} such that

$$r_{m*} = \min\{r_1, r_2, \dots, r_{\binom{n}{p}}\}, \tag{15.9}$$

so that $r_{m*} \leq r_m$, $m = 1, 2, \dots, \binom{n}{p}$. Teitz and Bart acknowledge that the p -median is not necessarily unique. They then addressed the task of finding it.

Teitz and Bart begin by outlining the direct enumeration method and Maranzana’s alternating algorithm, termed the *partition method* for obvious reasons. The former was deemed too computationally demanding and the latter of suspect robustness. This is followed by the main contribution, their vertex substitution method, which in their words “concentrates upon the formal definition of the generalized vertex median and its associated weighted distance matrix.”

The method proceeds as follows: for each possible subset of facility sites V_p^m , we may construct a submatrix \mathbf{R}_p^m by combining the relevant p columns as described above. Consider what happens when one vertex v_j in the facility subset is replaced by another vertex v_b outside this set; that is, the v_b column takes the place of the v_j column in \mathbf{R}_p^m . If r_{ij} is the i -th row minimum of \mathbf{R}_p^m , then its replacement by r_{ib} could have one of several outcomes:

If $r_{ib} \leq r_{ij}$, the increment to the i -th row contribution to sum r would be

$${}_i\Delta_{bj} = r_{ib} - r_{ij} \leq 0. \tag{15.10}$$

If $r_{ij} \leq r_{ib} \leq r_{is}$ (where r_{is} is the second-smallest i -th row element in \mathbf{R}_p^m),

$${}_i\Delta_{bj} = r_{ib} - r_{ij} \geq 0. \tag{15.11}$$

If $r_{ij} \leq r_{is} \leq r_{ib}$,

$${}_i\Delta_{bj} = r_{is} - r_{ij} \geq 0. \tag{15.12}$$

In the Teitz and Bart paper, the differences in these expressions are incorrectly reversed. For example, $r_{ib} - r_{ij}$ is incorrectly written $r_{ij} - r_{ib}$. The authors also seem to make a fundamental error by concluding that “if r_{ij} were not the i -th row minimum of \mathbf{R}_p^m , then no change in the i -th row contribution to r would result.” This is not generally true, as implied by the analysis above. There can be no increase in the objective value r , but v_b may still become the new closest facility to v_p , resulting in a reduction in r .

It is worth substituting vertex v_b for v_j only if the net effect of all increments

$$\Delta_{bj} = \sum_{i=1}^n {}_i\Delta_{bj} \tag{15.13}$$

is less than zero, i.e., if it reduces the total weighted distance. An iterative process of single vertex substitutions as suggested by Teitz and Bart may now be employed to obtain a monotone decreasing sequence of solutions that ends when a local minimum is reached.

Algorithm 3: Vertex Substitution (VS-1)

- Step 1:* Choose some initial facility subset V_1 containing p (randomly-selected) vertices.
- Step 2:* For each vertex $v_j \in V_1$, find its associated customer subset of vertices for which it is the closest facility (no rules are given for breaking ties as in Maranzana); compute the total weighted distance r_1 for the resulting solution.
- Step 3:* Select some vertex v_b not in the facility subset, i.e., $v_b \in V \setminus V_1$.
- Step 4:* Substitute v_b in turn for each vertex $v_j \in V_1$, and compute Δ_{bj} each time.
- Step 5:* Find that vertex $v_k \in V_1$ that, when replaced by v_b , most reduces the total weighted distance, that is,

$$\Delta_{bk} < 0 \text{ and } \Delta_{bk} = \min_j \{\Delta_{bj}\} \quad (15.14)$$

- Step 6:* If such a vertex v_k can be found, substitute v_b for v_k in the facility subset; label the new subset V_2 and compute $r_2 (= r_1 + \Delta_{bk})$. If no vertex v_k satisfies relation (15.14), simply retain the facility subset V_1 .
- Step 7:* Select another vertex, not previously tried, in the complement of V_1 , and repeat Steps 4 through 6.
- Step 8:* When all vertices in the complement of V_1 have been tried, define the resulting facility subset V_i as a new V_1 and repeat Steps 2 through 7. Call each such complete repetition a cycle.
- Step 9:* When one complete cycle results in no reduction in r , terminate the procedure. The output is the last solution obtained.

Note: It appears to be unclear in Step 7 whether the authors intended that the new subset V_2 replace the original subset V_1 in the repetition of Steps 4–6. However, this would only affect the type of improvement strategy utilized, and not the gist of the procedure. We interpret the authors' intention as using the original V_1 in each such repetition of Steps 4–6, giving rise to what is termed today a “best” improvement strategy; i.e., look at all solutions in the one-interchange neighborhood of V_1 and select the best one. It is also unclear how they intend us to perform Step 6 in successive iterations; their use of V_i in Step 8 suggests that they would label further subsets V_3, V_4, \dots, V_i . This is not necessary, since we need only maintain a current “best” substitution at this step, which we could label V_i throughout. Moreover, the total weighted distance need not be calculated each time.

Teitz and Bart acknowledge that a situation could arise in which a single vertex substitution produces no further improvement, whereas pairwise or higher substitutions would further reduce the total weighted distance. However, they do not characterize this case or give examples. In their experiments on random graphs with 25 vertices, they observe that the procedure always terminated (i.e., reached a

local minimum) within four cycles. Most important—they note a very significant improvement in solution quality using vertex substitution compared to the partition method. The partition method furthermore exhibits considerable variation in performance. They conclude that vertex substitution is the preferable heuristic.

15.3 Impact of the Early Heuristics

This section investigates what further developments have been made on the basis of the heuristic methods described in the previous section. We first examine work that considers generalization of distance measurements, followed by a variety of location—allocation models and modern heuristics in continuous and discrete spaces.

15.3.1 *Generalization of Distance Measurements*

The Maranzana and Teitz and Bart approaches were defined in network space, but it is not necessary to have a network structure to define the p -median or to solve it using their procedures. Many examples in the literature do not rely on an underlying network structure. Both discrete and network spaces are defined with reference to a matrix of shortest distances, and operate through consideration of these internodal distances. This is possible because the Hakimi (1964) finding ensures that optimal locations can be limited to the vertices—hence the internodal distance matrix is all the information required. It follows that the relevant distance matrix among pairs of “vertices” can be defined other than within a network or if, within a network, without specifying the network structure. We can apply the partition and vertex substitution methods to any system where a matrix $[D_{ij}]$ is provided. Given a set of nodes in space, these distances might be specified for example as Euclidean distances, airline travel times, psychologically perceived travel costs, or in many other different ways.

The Teitz and Bart and Maranzana papers work on the assumption that the customer set represents the potential facility locations from which trial sites can be selected. In modern practice it is recognized that this is often not realistic. Some customer sites may be unsuited to facilities; some ideal facility locations may not express demand. Thus, we recommend that a separate set of potential facility locations be maintained in working with network and discrete space models. The distance matrix would be constructed and used in a similar fashion as before. In some realistic problems, facilities may already exist in some locations, and the heuristics above are easily adapted to deal with this situation.

The Cooper papers (1963, 1964) assumed that distances are measured by the Euclidean norm. Since that time more general distance functions, such as the l_p norm, have been incorporated in location models to provide more accurate measures of

travel distance. Given any two points, $X_1=(x_1, y_1)$, $X_2=(x_2, y_2)$ in the plane, the ℓ_p distance between them is given by:

$$\ell_p(X_1, X_2) = [|x_1 - x_2|^p + |y_1 - y_2|^p]^{1/p}$$

where the parameter $p \geq 1$. When $p=1$, we have the well-known rectangular (or Manhattan) norm; the Euclidean norm occurs with $p=2$. The Cooper algorithm is readily extended to the median problem with ℓ_p distances after modifying the Weiszfeld formulas appropriately; see, e.g., Love et al. (1988). The use of “block norms” allows the location step to be solved by linear programming techniques. In fact the problem may now be reduced from continuous space to a finite set of intersection points, thus allowing a vertex substitution heuristic to be used as well.

The distance function may be raised to some power in order to model more effectively the transportation costs or times; an example is the fire engine travel time study in Kolesar et al. (1975). Geodesic distances are typically used for location on a sphere as in the case of air travel. In *cluster analysis*, which has important applications, for example, in data mining, the fixed points (vertices) are located in higher-dimensional space according to the number of attributes involved. The well-known k -means model from the data mining literature is simply a version of the p -median model ($p=k$) with squared Euclidean distances.

The partition and vertex substitution methods have been readily adapted to such generalizations of the original continuous and discrete (network) problems.

15.3.2 Other Location-Allocation Models

The general principles of partition and vertex substitution may be extended to other forms of the location-allocation problem as introduced elsewhere in this book. With the vertex substitution method, we simply revise the procedure for calculating the incremental change in objective function associated with each swap move of vertex entering and vertex leaving the solution. The generation of a monotone sequence and convergence to a ‘local’ optimum are guaranteed. With the partition method, the location step is adjusted according to the type of objective function under consideration. For problems that are *separable* into location and allocation phases, we may show again that the sequence generated is monotone, which is essential for convergence of the heuristic. However, the partition method may not converge in more general cases. Consider, as an example, a form of the *covering* problem where the goal is to locate sensors on a grid in order to maximize the mean probability of detection measured at the grid points. The location step is no longer separable into p independent single facility problems due to additional interactions that exist with facilities (sensors) other than the closest one, and thus the partition method breaks down. A similar situation may occur in location-allocation models involving *noxious* facilities. It therefore appears that the partition method is not as universally

useful as the vertex substitution method for problems occurring on networks or in discrete space.

15.3.3 *Modern-Day Heuristics*

The partition and vertex substitution methods both fall in the category of *local search*; that is, the procedure finds a better solution in a local neighborhood of the current solution and iterates in this fashion until a local optimum is reached. It is interesting to note that in network (or discrete) space, any local optimum obtained by the vertex substitution method must also be a local optimum in the partition method, but the converse is not necessarily true. This is due to the fact that the location step in the partition method is equivalent to a “restricted” set of vertex swap moves. Thus, starting from some local optimum, a better partition of the customer set may be found by examining all possible swap moves as in the vertex substitution method. This may explain the superiority observed by Teitz and Bart of their heuristic, as well as the higher variability of results obtained by the partition method. It is also interesting to note that comparative studies of the two methods appear to be limited to the experiment of Teitz and Bart on a few small random instances, and some further testing by Rosing et al. (1979). Yet the vertex substitution method is widely used to this day, while Maranzana’s work in comparison has been largely ignored. There may be computational advantages, for example, in using a two-stage approach where the fast partition method is applied first on a random initial solution, followed by vertex substitution with the solution from the first stage as the starting point.

The Cooper algorithm is still widely used on problems posed in continuous space. A few variants that seem to work better have been suggested including, as noted, Scott (1971) who starts with an initial random set of facility locations instead of an initial allocation. Care must be taken, since the Cooper algorithm may lead to a degenerate solution (Brimberg and Mladenovic 1999) where some facilities end up having no customers assigned to them. The shortcoming is easily remedied by inserting such facilities at unoccupied vertex locations (those customers that do not have coinciding facilities) whenever the situation arises within the solution process.

As problem size defined by n and p increases, an exponential growth in the number of local optima is observed. Thus, local search methods become inefficient. We will see next that the partition and vertex substitution methods still play an important role in the more advanced techniques used today.

15.3.3.1 **The Continuous p -Median Problem**

The random multi-start version of Cooper’s algorithm remained the state-of-the-art for many years despite a number of other competing heuristics. Notable among these is a heuristic developed by Love and Juel (1982) that is the first method to

impose a set of neighborhood structures on the problem. A given neighborhood of a solution is defined here as the set of points around that solution that are obtained by exchanging a specified number of assignments of customers from their current facilities to new ones. The authors consider up to two exchanges, and show that the two-exchange neighborhood may be used (at a computational cost, of course) to ‘jump out’ of a local optimum trap in the one-exchange neighborhood. In their procedure the facilities are always optimally located with respect to any given allocation of the customers. Other heuristics include gradient-based methods (Murtagh and Niwattisyawong 1982, and Chen 1983) and a projection method by Bongartz et al. (1994). For further details, see, for example, the survey paper by Brimberg et al. (2008a).

Recall that one of Cooper’s initial ideas was to solve a discrete version of the problem where the facility locations are restricted to the set of fixed points given by the customers and the shortest-path distance is simply the Euclidean distance between each pair of vertices. Hansen et al. (1998) revisit this idea several years later while taking advantage of an efficient code by Hanjoul and Peeters (1985) to solve the discrete problem exactly. A second stage involves a continuous improvement where p single facility problems resulting from the partition of the customer set by the discrete solution are solved. Excellent results are reported, but computation times become excessive. Brimberg et al. (2000) propose a new neighborhood structure based on the vertex substitution idea of Teitz and Bart (1968); that is, facilities are relocated one at a time to an unoccupied fixed point (a customer that does not have a coincident facility). The one-interchange neighborhood contains all such possible single moves. A local search using Cooper’s algorithm is then conducted from all or selected points in this neighborhood. The authors investigate various “drop and add” strategies in the selection process, which allow a reduction in the size of the neighborhood from $O(np)$ to $O(n+p)$, and as a result, a much faster local search. When the full one-interchange neighborhood is verified, an efficient updating procedure by Whitaker (1983) is used. The relocation heuristics are able to obtain better results than the multi-start Cooper algorithm in a fraction of the time.

The recent application of *metaheuristics* to the continuous p -median problem has resulted in a significant advance in the state-of-the-art. Unlike local search that examines a narrow region of the solution space and terminates at a local optimum, metaheuristics are general frameworks that allow the search to expand to different regions of the solution space, and thus escape the “local optimum trap.” A comparative study (Brimberg et al. 2000) shows that as problem size increases (and the number of local minima explodes), the performance of the multi-start Cooper algorithm deteriorates significantly relative to new heuristics based on Tabu search, variable neighborhood search, and the genetic algorithm. It is interesting to note, however, that these newer methods usually have Cooper’s algorithm embedded within them. For example, the various versions of variable neighborhood search in the above comparative study use Cooper’s algorithm in the local search step. The initial population in the genetic algorithm of Houck et al. (1996) is obtained by repeating Cooper’s algorithm from random starting points until an adequate number of local minima is found, and after the crossover

operation, the new solution is improved (mutation operation) using the Cooper algorithm. For a further update on metaheuristic-based methods for solving the continuous p -median problem, see Brimberg et al. (2008a).

15.3.3.2 The Discrete p -Median Problem

As noted in Mladenovic et al. (2007), the vertex substitution method by Teitz and Bart, which they refer to as the *Interchange* procedure, is still “commonly used as a standard to compare with other methods.” Both Maranzana’s partition method, aptly named the *Alternate* heuristic, and the Interchange procedure have been used in composite type heuristics. For example, Captivo (1991) adds facilities one at a time in a greedy fashion that reduces total cost as much as possible, and then uses the Alternate procedure in each step to further improve the solution. Another composite method first constructs a greedy solution and then applies the Interchange procedure to that solution; it is often used for comparison with other new methods (see Voss 1996, and Hansen and Mladenovic 1997). Lagrangian-based procedures that “alternate” between solving for the primal variables and adjusting the Lagrange multipliers typically use a local search as above to improve the obtained solution (as in Beasley 1993).

The concept of neighborhood structure is intimately related to the vertex substitution method. We may view this method as a local search in the *1-interchange* neighborhood. Generalizations are now possible. For example, Kochetov et al. (2005) propose a new neighborhood structure, termed *LK* (Lin-Kernighan), which employs a depth parameter k that counts the number of interchange moves within one step of local search. The *LK*(k) neighborhood may be described as follows: (i) find two vertices v_{add} and v_{drop} that give the best solution in the 1-interchange neighborhood; (ii) exchange these two vertices to get a new solution; (iii) repeat the above steps k times, not allowing any facility that has been dropped to re-enter the solution. The process is repeated until a local minimum in the *LK* neighborhood is reached. This type of local search has been used within Lagrangian relaxation, random rounding (after linear relaxation), and ant colony optimization (Dorigo and Di Caro 1999). The 1-interchange neighborhood can be modified in a straightforward way to handle the related *simple plant location problem*. In this case a fixed cost f_i is charged to open a facility at vertex v_p , and the number of facilities to open is unknown. Brimberg et al. (2008b) examine an extended version of the simple plant location problem with nonlinear objective function representing the return on investment. They use an expanded local search neighborhood that allows all single moves where either a vertex is opened (v_{add}), a vertex is closed (v_{drop}), or an interchange is made (v_{add} and v_{drop}).

Mladenovic et al. (2007) note that “the Interchange method is one of the most often used classical heuristics either alone or as a subroutine of other more complex methods or within metaheuristics.” In large scale applications it is therefore critical that the procedure be implemented in an efficient manner. The popular CLARANS (Clustering Large Applications based on RANdomized Search) method in data min-

ing (Ng and Han 2002) conducts a local search using a small random sample of points in the 1-interchange neighborhood. Efficient implementations that can evaluate in reasonable time the entire neighborhood for very large instances, including the fast interchange of Whitaker (1983) mentioned previously, are summarized in Mladenovic et al. (2007).

Several modern heuristics that derive from metaheuristic rules use the vertex substitution method in some form. In the *Tabu search* procedure of Mladenovic et al. (1996), the 1-interchange move is extended to what they term the 1-chain-substitution move. In Rolland et al. (1996), the 1-interchange move is divided into add and drop moves that do not necessarily follow each other in an approach within Tabu search known as *strategic oscillation*. Note that this procedure allows the trajectory that is generated to oscillate between feasible and infeasible solutions. Kochetov (2001) proposes a simple probabilistic Tabu search in which a restricted (random) 1-interchange neighborhood is used. The *simulated annealing* heuristic of Chiyoshi and Galvao (2000) combines the 1-interchange neighborhood with the general methodology of simulated annealing. The *scatter search* method of Garcia-Lopez et al. (2003) uses the Interchange procedure in a final step to improve the combined solutions that are obtained.

The *Variable Neighborhood Search* methodology imposes a set of neighborhood structures on the solution space in order to conduct a systematic search at different distances from the current solution. The movement to different neighborhoods is accomplished by a ‘shaking’ operation. In the standard approach for the p -median problem (e.g., Hansen and Mladenovic 1997), the neighborhood structures are defined by moving 1, 2, ..., k_{max} facilities from their currently occupied vertices to new unoccupied ones. The shaking operator thus selects a random point in the k -neighborhood by randomly moving k facilities in this manner. A local search from this point is conducted using the Interchange procedure.

Heuristic concentration (Rosing and ReVelle 1997) is a metaheuristic of special interest to this chapter as it was developed specifically for the p -median problem and is based straightforwardly on the Teitz and Bart algorithm. The local minima arising in repeated runs of the Teitz and Bart algorithm are identified as two-, three-, and so on “traps” by Rosing and Hodgson (2002); these identify clusters of nodes that cannot be avoided by single interchanges. The main idea behind heuristic concentration is to then create a concentration set of desirable facility sites (open facility sites that most often appeared in the solutions from the first stage), and use this concentration set as the set of potential facility locations, thus reducing the solution space of the problem. Nodes that occur in all solutions may be assumed to be in the optimal solution if so desired. The much smaller problem defined by the demand nodes and the concentration set may be solved optimally or approximately. Heuristic concentration has been shown to provide very good, usually optimal, solutions to problems of several hundred nodes, although these are considered relatively small instances by today’s standards. Since the number of customers remains at its original size, computation times may become unmanageable for larger problems. This shortcoming may be addressed in the future by applying neighborhood approaches to the traps identified in the concentration set.

Thus we see that several of the new methods have at their hearts the fundamental interchange notion of Teitz and Bart (1968). Mladenovic et al. (2007) conclude that the more recent heuristics outperform Teitz and Bart, but due to the number of different data structures and implementations, they are unable to conclude “what metaheuristic dominates others.” Another useful source to note is Reese (2006).

15.4 Future Research

We believe that interest in heuristics will continue to grow in the coming years in studies of combinatorial and global optimization problems including, of course, the p -median problem. Here are some reasons why.

Networks are getting increasingly larger in real applications. For example, Brimberg et al. (2000) motivate their work by citing two actual case studies, a transshipment center location problem and a districting problem with $(p, n) = (20, 1,700)$ and $(170, 1,400)$, respectively. A very large-scale study dealing with spare parts logistics for a Japanese manufacturing company with 6,000 customers and 380,000 potential warehouse sites is cited in Brimberg et al. (2008a). With such trends as globalization of business, we can expect the size and complexity of location and distribution problems to only increase. Given the limiting assumptions inherent in mathematical models, finding the “optimal” solution, even if it were possible, may be of questionable importance in practice. It seems a more sensible approach would be to use heuristics to find a set of alternative “good” solutions in a way that strikes a proper balance between quality of solution and computing time.

New important applications of the p -median model are materializing that are outside the original scope of locating physical assets such as warehouses. We mention as an example the importance of the p -median model and other related models in the field of *data mining*. One objective in data mining is to detect useful patterns within databases by using models such as the p -median that are able to partition the dataset into meaningful clusters. These databases generally contain several thousand entries, and thus the use of heuristics becomes a practical necessity.

New developments within the field of heuristics are extending the usefulness of these methods. For example, using decomposition to solve a series of smaller (decomposed) problems is proving to be a highly effective and efficient approach to handle large problem instances. In Hansen et al. (2001), a decomposition variant of variable neighborhood search, referred to as variable neighborhood decomposition search, obtains notably better results than basic variable neighborhood search in less computing time. In fact, the method finds much better results than fast-interchange in the same time fast-interchange takes for a single descent. Another example is the recent development of primal-dual heuristics that are able to obtain tight lower bounds on the optimal solution of the p -median and related simple plant location problems by solving exactly or approximately a relaxed version of the dual, see Hansen et al. (2007). Thus, a guaranteed bound on the quality of the solution obtained by the primal heuristic is now provided. Heuristics are also used in con-

junction with exact solution methods. Since these exact methods are generally very sensitive to the starting point, it is important to use a good heuristic initially. Brimberg et al. (2000) note that the improved solution quality from the newer heuristics available has enabled the exact solution of much larger instances of the continuous p -median problem than before.

We have seen a tremendous growth in the field of heuristics in recent years. Empirical studies have shown consistently that the new metaheuristics at our disposal work better than the older methods, but aside from this we do not understand much of what transpires. Research into the theoretical underpinnings of metaheuristics is still in its infancy, but judging from the recent *Seventh Metaheuristics International Conference* (June 2007), this is becoming a very hot area indeed. It seems a safe bet to predict that the field of metaheuristics will be subject to rigorous theoretical analysis in the years to come. We believe that statistical studies of the *landscapes* derived from various local search operators will play a useful role here. The ultimate goal, aside from designing better heuristics, will be a deeper understanding of the fundamental nature of combinatorial and global optimization problems.

15.5 Conclusions

We have reviewed the classical heuristics introduced by Cooper (1963, 1964), Maranzana (1964), and Teitz and Bart (1968). Maranzana's paper, although important also for its formulation of the network model and some fundamental results, may be considered less significant as its partition method was quickly superseded by the vertex substitution method of Teitz and Bart. The contributions of these original papers were timely and important, as they introduced a wide audience to some fundamental location problems and showed how they could be solved. They stand out as important way posts in location science and are truly deserving of their celebrity status.

Rather than review derivative work in detail, we have guided readers to the detailed reviews by Brimberg et al. (2000), Reese (2006), Mladenovic et al. (2007), and Brimberg et al. (2008a). These reviews indicate that the performance of the classical heuristics suffers in the face of the explosive number of local minima that arise in large problems. None of these classic approaches is ready for retirement, however.

The Cooper algorithm, accompanied by graphical illustration, is an excellent tool for teaching a fundamental lesson of location modeling and optimization in the classroom, as aptly demonstrated in Scott (1971).

The vertex substitution method of Teitz and Bart lies at the heart of all the interchange-based heuristics developed to solve the p -median problem more effectively. Although much progress has been made in recent years in the development of metaheuristics, much work still remains to understand the underlying theory, and in consequence, to design heuristics in a more intelligent fashion.

References

- Beasley JE (1993) Lagrangian heuristics for location problems. *Eur J Oper Res* 65:383–399
- Bellman R (1958) On a routing problem. *Q Appl Math* 16:87–90
- Bongartz I, Calamai PH, Conn AR (1994) A projection method for ℓ_p norm location-allocation problems. *Math Program* 66:283–312
- Brimberg J, Mladenovic N (1999) Degeneracy in the multi-source Weber problem. *Math Program* 85:213–220
- Brimberg J, Hansen P, Mladenovic N, Taillard ED (2000) Improvements and comparison of heuristics for solving the uncapacitated multisource Weber problem. *Oper Res* 48:444–460
- Brimberg J, Hansen P, Mladenovic N, Salhi S (2008a) A survey of solution methods for the continuous location-allocation problem. *Int J Oper Res* 5:1–12
- Brimberg J, Hansen P, Laporte G, Mladenovic N, Urosevic D (2008b) The maximum return-on-investment plant location problem with market share. *J Oper Res Soc* 59:399–406
- Captivo EM (1991) Fast primal and dual heuristics for the p -median location problem. *Eur J Oper Res* 52:65–74
- Chen R (1983) Solution of minisum and minimax location-allocation problems with Euclidean distances. *Nav Res Logist Q* 30:449–459
- Chiyoshi F, Galvao RD (2000) A statistical analysis of simulated annealing applied to the p -median problem. *Ann Oper Res* 96:61–74
- Cooper L (1963) Location-allocation problems. *Oper Res* 11:331–343
- Cooper L (1964) Heuristic methods for location-allocation problems. *SIAM Rev* 6:37–53
- Dorigo M, Di Caro G (1999) The ant colony optimization meta-heuristic. In: Corne G, Dorigo M, Glover F (eds) *New ideas in optimization*, McGraw-Hill, New York
- Garcia-Lopez F, Melian Batista B, Moreno Perez JA, Moreno Vega JM (2003) Parallelization of the scatter search for the p -median problem. *Parallel Comput* 29:575–589
- Hakimi SL (1964) Optimum locations of switching centers and the absolute centers and medians of a graph. *Oper Res* 12:450–459
- Hanjoul P, Peeters D (1985) A comparison of two dual-based procedures for solving the p -median problem. *Eur J Oper Res* 20:387–396
- Hansen P, Mladenovic N (1997) Variable neighborhood search for the p -median. *Locat Sci* 5:207–226
- Hansen P, Mladenovic N, Taillard E (1998) Heuristic solution of the multisource Weber problem as a p -median problem. *Oper Res Lett* 22:55–62
- Hansen P, Mladenovic N, Perez-Brito D (2001) Variable neighborhood decomposition search. *J Heuristics* 7:335–350
- Hansen P, Brimberg J, Urosevic D, Mladenovic N (2007) Primal-dual variable neighborhood search for the simple plant-location problem. *INFORMS J Comput* 19:552–564
- Houck CR, Joines JA, Kay MG (1996) Comparison of genetic algorithms, random restart and two-opt switching for solving large location-allocation problems. *Comput Oper Res* 23:587–596
- Kariv O, Hakimi SL (1979) An algorithmic approach to network location problems: The p medians. *SIAM J Appl Math* 37:539–560
- Katz IN (1974) Local convergence in Fermat's problem. *Math Program* 6:89–104
- Kochetov Y (2001) Probabilistic local search algorithms for the discrete optimization problems. *Discrete Mathematics and Applications*. Moscow, MSU, 84–117 (in Russian)
- Kochetov Y, Alekseeva E, Levanova T, Loresh N (2005) Large neighborhood search for the p -median problem. *Yugosl J Oper Res* 15:53–63
- Kolesar P, Walker WE, Hausner J (1975) Determining the relation between fire engine travel times and travel distances in New York City. *Oper Res* 23:614–627
- Kuhn HW (1973) A note on Fermat's problem. *Math Program* 4:98–107
- Kuhn HW, Kuenne RE (1962) An efficient algorithm for the numerical solution of the generalized Weber problem in spatial economics. *J Reg Sci* 4:21–34

- Love RF, Juel H (1982) Properties and solution methods for large location-allocation problems. *J Oper Res Soc* 33:443–452
- Love RF, Morris JG, Wesolowsky GO (1988) *Facilities location: models and methods*. North Holland, New York
- Maranzana FE (1964) On the location of supply points to minimize transport costs. *Oper Res Q* 15:261–270
- Megiddo M, Supowit KJ (1984) On the complexity of some common geometric location problems. *SIAM J Comput* 13:182–196
- Mladenovic N, Moreno-Perez JA, Moreno-Vega JM (1996) A chain-interchange heuristic method. *Yugosl J Oper Res* 6:41–54
- Mladenovic N, Brimberg J, Hansen P, Moreno-Perez JA (2007) The p -median problem: A survey of metaheuristic approaches. *Eur J Oper Res* 179:927–939
- Murtagh BA, Niwattisyawong SR (1982) An efficient method for the multi-depot location-allocation problem. *J Oper Res Soc* 33:629–634
- Ng R, Han J (2002) CLARANS: a method for clustering objects for spatial data mining. *IEEE Trans Knowl Data Eng* 14:1003–1016
- Reese J (2006) Solution methods for the p -median problem: an annotated bibliography. *Networks* 48:125–142
- Rolland E, Schilling DA, Current JA (1996) An efficient tabu search procedure for the p -median problem. *Eur J Oper Res* 96:329–342
- Rosing KE, Hodgson MJ (2002) Heuristic concentration for the p -median: An example demonstrating how and why it works. *Comput Oper Res* 29:1317–1330
- Rosing KE, ReVelle CS (1997) Heuristic concentration: Two stage solution construction. *Eur J Oper Res* 96:329–342
- Rosing KE, Hillsman EL, Rosing-Vogelaar H (1979) A note comparing optimal and heuristic solutions to the p -median problem. *Geogr Anal* 11:86–89
- Scott AJ (1971) *Combinatorial Programming, spatial analysis and planning*. Methuen, London
- Teitz MB, Bart P (1968) Heuristic methods for estimating the generalized vertex median of a weighted graph. *Oper Res* 16:955–961
- Voss S (1996) A reverse elimination approach for the p -median problem. *Stud Locat Anal* 8:49–58
- Weiszfeld E (1937) Sur le point pour lequel la somme des distances de n points donnés est minimum. *Tohoku Math J* 43:355–386
- Whitaker R (1983) A fast algorithm for the greedy-interchange for large-scale clustering and median location problems. *Inf Syst Oper Res* 21:95–108