

Optimizing Emergency Ambulance Operations Using Operations Research Techniques

Author Details:

Prapti Lakhani¹, Raghav Agarwal², Pragun Bindra³, Palash Sethia⁴, Palak Doshi⁵

¹ BSc Finance, Anil Surendra Modi School of Commerce (ASMSOC), SVKM's NMIMS, Mumbai, Maharashtra, India

Corresponding author email: lakhaniprapti2007@gmail.com



<https://doi.org/10.55041/ijst.v2i10.040>

Cite this Article: Lakhani, P., Agarwal, R., Bindra, P., Sethia, P. & Doshi, P. (2026). Optimizing Emergency Ambulance Operations Using Operations Research Techniques. *International Journal of Science, Strategic Management and Technology*, 02(10), 1-9. <https://doi.org/10.55041/ijst.v2i10.040>

License:  This article is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited.

Abstract

Emergency medical services must decide where to put ambulances how many of each type to keep, which ambulance to send to which call and which road to choose in just a few minutes. Emergency medical services can benefit from operations research tools taught at university level and this paper shows how these tools can help Emergency medical services make all of these decisions in a single linked system. Using a sample data set for a made-up city district with eight demand zones and five possible stations a linear programming problem fixes the mix of Advanced and Basic Life Support ambulances; a set-covering integer program chooses stations that can reach every zone in eight minutes; a transportation model solved by Vogel's Approximation Method and verified with the MODI method assigns trips from stations to demand clusters; the Hungarian method matches free ambulances to calls that happen at the same time; Dijkstra's algorithm finds the quickest road path; and a nearest-neighbour heuristic, a 2-opt improvement and simulated annealing design a multi-stop patient transport route. The linear programming solution gives 3 ALS and 8 BLS units that serve 82 trips each day; three of the five stations provide eight-minute coverage; VAM achieves an optimal transportation cost of 377 minutes, which is 6.7 percent lower than the North-West Corner starting point; the Hungarian assignment reduces total response time from 35 to 30 minutes, a 14.3 percent improvement compared with dispatching the nearest available ambulance; and both 2-opt and simulated annealing cut the route time from 77.4 to 60.0 minutes, a 22.5 percent reduction that matches the exact optimum. The results show that simple spreadsheet-ready operations research tools give measurable benefits, for Emergency medical services planning.

Keywords: ambulance routing; emergency medical services; Hungarian method; transportation problem; set covering; simulated annealing

1. Introduction

In an emergency every single minute between the call and the arrival of the ambulance matters a lot for the patient's chance of survival. For out-of-hospital arrest the chance of survival drops by about 5.5% for each minute that CPR, defibrillation and advanced care are delayed (Larsen et al., 1993). That is why response time is the important performance measure for any emergency medical service or EMS.

Response time is not about how fast a driver can go. It depends on a series of planning decisions. Where are the ambulances based? How many of each type are available? Which free ambulance is sent to which call?. Which route does it take? Each of these is a resource-allocation problem.. Each one has a standard model in operations research or OR.

The challenge becomes even greater in cities. Traffic is dense. Roads are uneven.. There is a mix of public services (like 108) and private ambulance operators. Simple dispatch rules—like sending the available vehicle—are easy to use.. They do not always give the best result. Especially when multiple calls come in at once or when the station locations are poorly chosen.

Most research on ambulance planning uses stochastic models and simulations (Aringhieri et al. 2017; Bélanger et al., 2019). These models work well. They are hard for students or small operators to use. This paper takes an approach. It shows that the basic OR techniques taught in a course can already lead to clear measurable improvements—when applied step by step to one EMS system.

The objectives of this study are:

- 1.To model the EMS planning decisions—fleet mix, station siting, trip allocation, dispatch and routing—using standard operations research methods.
- 2.To solve each model in a case study using linear programming, integer programming, the transportation method (VAM and MODI) the Hungarian assignment method, Dijkstra's shortest-path algorithm, a nearest-neighbour heuristic with 2-opt improvement and simulated annealing.
- 3.To compare optimized decisions with rules of thumb and measure the improvement in response time.The rest of the paper is organized as follows. Section 2 reviews the literature. Section 3 describes the case setting and the six models. Section 4. Discusses the results. Section 5 concludes with limitations and directions, for work.

2. Literature Review

OR research on ambulances falls into three streams: where to place them, how to dispatch them, and how to route them. Each stream builds on a classical model that this paper reuses.

2.1 Location and coverage

The location set covering problem (LSCP) was put forward by Toregas et al. (1971) to determine the minimum number of stations required to service all demand points within a given time frame. Church and ReVelle (1974) took a more lenient approach with the maximal covering location problem, using a set number of vehicles to cover the greatest amount of demand. Later work has seen Daskin (1983) factor in the possibility of the closest ambulance being occupied, and Gendreau et al. (1997) employ tabu search on a double-coverage variant. For an account of how such models have developed into those for relocation and dynamic deployment, one can refer to the reviews by Brotcorne et al. (2003) and Bélanger et al. (2019).

2.2 Allocation and assignment

Hitchcock (1941) was the first to present the transportation problem as a means of moving units from source to destination at the lowest cost. The North-West Corner rule, Vogel's Approximation Method and the MODI test are still the way the subject is taught (Taha, 2017; Hillier & Lieberman, 2015). In the case of the assignment problem, Kuhn's (1955) Hungarian method provides an efficient solution. Within EMS these are well suited to the task of assigning ambulances to calls or stations to areas of need.

2.3 Routing

For finding the quickest route on a road network there is Dijkstra's (1959) standard algorithm. But if a single vehicle has to make multiple stops the matter is one of travelling salesman or vehicle routing (Dantzig & Ramser, 1959), and exact solutions become difficult with scale. Heuristics like the savings method (Clarke & Wright, 1964) or nearest-

neighbour construction will produce a route in short order. Where better results are called for, metaheuristics are applied, be it simulated annealing (Kirkpatrick et al., 1983), tabu search (Glover, 1986), genetic algorithms (Holland, 1975) or ant colony optimisation (Dorigo & Gambardella, 1997).

2.4 Research gap. Reviews by Brotcorne et al. (2003), Aringhieri et al. (2017) and Bélanger et al. (2019) show that most EMS studies use one advanced model at a time. Few papers show, at an introductory level, how the basic toolkit works together across all planning stages on one system. This paper fills that gap.

Study	Technique	Main contribution
Toregas et al. (1971)	Set covering (ILP)	Minimum stations for full coverage
Church & ReVelle (1974)	Maximal covering	Most demand covered with fixed fleet
Daskin (1983)	Expected covering	Accounts for busy ambulances
Kuhn (1955)	Hungarian method	Optimal one-to-one assignment
Dijkstra (1959)	Shortest path	Fastest route on a network
Kirkpatrick et al. (1983)	Simulated annealing	Escapes local optima in routing
Gendreau et al. (1997)	Tabu search	Double coverage ambulance location
Bélanger et al. (2019)	Review	Location, relocation and dispatch trends

3. Methodology

The study uses a case study approach. Six OR models are used one after another in a single example EMS district. Each model answers one planning question (see Figure 1). The data is not real. It was made up by the authors to show how each method works. The values used are, in ranges that make sense for a city setting. These data are not matched to any EMS system. They are kept simple so each step can be checked by hand. All models were also solved in Python using PuLP with the CBC solver (Mitchell et al., 2011) SciPy (Virtanen et al. 2020) and NetworkX (Hagberg et al., 2008) to verify the results. Figures were drawn with Matplotlib (Hunter, 2007).




Figure 1. Research framework. The output of each planning model feeds the next; the operational models then guide each trip.

3.1 Case setting. The district has eight demand zones (Z1 to Z8) and five candidate ambulance stations (S1 to S5). The response-time standard is 8 minutes, following the 480-second travel-time benchmark for advanced life support units in

NFPA 1710 (National Fire Protection Association [NFPA], 2020). Table 1 gives the estimated travel time from each candidate station to each zone.

Table 1. Travel time from candidate station to demand zone (minutes)

Station	Z1	Z2	Z3	Z4	Z5	Z6	Z7	Z8
S1	4	7	12	15	9	14	18	11
S2	9	5	6	8	13	9	14	16
S3	14	11	7	5	16	8	9	18
S4	6	12	15	17	5	13	16	7
S5	12	9	11	9	10	6	7	13

3.2 Model 1: Fleet mix (LPP). Let x_1 be the number of Advanced Life Support (ALS) ambulances and x_2 the number of Basic Life Support (BLS) ambulances. An ALS unit completes 6 trips a day and a BLS unit 8 (ALS units spend longer on scene with critical patients). Daily operating cost is ₹4,000 per ALS and ₹3,000 per BLS, against a budget of ₹36,000. An ALS unit needs 3 crew and a BLS unit 2, from a pool of 26 staff. At least 3 ALS units are required for critical cases.

Maximise $Z = 6x_1 + 8x_2$

subject to $4x_1 + 3x_2 \leq 36$; $3x_1 + 2x_2 \leq 26$; $x_1 \geq 3$; $x_1, x_2 \geq 0$

The model is solved by the graphical method, checking each corner point of the feasible region.

3.3 Model 2: Station siting (set covering ILP). Let $y_i = 1$ if station i is opened and 0 otherwise, and let $a_{ij} = 1$ if station i reaches zone j within 8 minutes (from Table 1).

Minimise $\sum_i y_i$ ($i = 1, \dots, 5$)

subject to $\sum_i a_{ij} y_i \geq 1$ for every zone $j = 1, \dots, 8$; $y_i \in \{0, 1\}$

3.4 Model 3: Trip allocation (transportation problem). The opened stations supply daily trip capacity to four demand clusters (D1 to D4). The cost c_{ij} is the average response time in minutes. Supply and demand are balanced at 60 trips per day (Table 2). The initial solution is found by the North-West Corner (NWC) rule and by VAM, and the MODI method tests optimality using $u_i + v_j = c_{ij}$ for occupied cells and the opportunity cost $d_{ij} = c_{ij} - (u_i + v_j)$ for empty cells. The solution is optimal when every $d_{ij} \geq 0$.

Table 2. Transportation data: response time (minutes), supply and demand (trips per day)

Station	D1	D2	D3	D4	Supply
S4	6	9	11	7	20
S2	8	5	10	12	15
S5	12	10	6	8	25
Demand	18	12	16	14	60

3.5 Model 4: Dispatch (assignment problem). At peak time, four calls (C1 to C4) arrive together and four ambulances (A1 to A4) are free. Table 3 gives the time each ambulance needs to reach each call. The Hungarian method (row reduction, column reduction, covering zeros with the fewest lines, and adjusting by the smallest uncovered value) gives the one-to-one assignment with least total time. It is compared with the common rule of sending the nearest free ambulance to each call in order of arrival.

Table 3. Time for ambulance to reach call (minutes)

Ambulance	C1	C2	C3	C4
A1	12	12	16	5
A2	11	10	13	14
A3	9	13	6	6
A4	17	8	6	6

3.6 Model 5: Fastest path (Dijkstra's algorithm). Once dispatched, an ambulance at station S must reach hospital H through the road network in Figure 2. Each edge carries a travel time in minutes. Dijkstra's algorithm labels nodes in order of their shortest known time from S until H is fixed.

Fastest path $S \rightarrow B \rightarrow A \rightarrow C \rightarrow E \rightarrow H = 15$ min (edge = minutes; [] = Dijkstra label)




Figure 2. Road network from station S to hospital H, with travel times in minutes. The highlighted path is the result of Model 5 (Section 4.5).

3.7 Model 6: Multi-stop route (heuristic and metaheuristic). A BLS unit based at depot D must collect nine non-emergency patients (P1 to P9, for example dialysis or discharge transfers) and return. Point coordinates are in kilometres and travel time is taken as 1.5 minutes per kilometre of straight-line distance. Three methods are compared: (a) the nearest-neighbour heuristic, which always visits the closest unvisited point (Rosenkrantz et al., 1977); (b) 2-opt improvement, which reverses route segments while this shortens the tour (Croes, 1958); and (c) simulated annealing, which also accepts some worse moves early on with probability $e^{(-\Delta/T)}$, starting at $T = 50$ and cooling by a factor of 0.9995 over 20,000 iterations. Because the instance is small, all 362,880 possible routes were also enumerated to find the exact optimum as a benchmark.

Table 4. Pickup point coordinates (km)

Point	D	P1	P2	P3	P4	P5	P6	P7	P8	P9
x	0	2	5	8	9	6	3	11	1	7
y	0	6	9	7	3	1	3	9	10	5

4. Results and Discussion

Every optimised decision matched or beat the rule of thumb it replaced, with gains of 6.7% to 22.5% in response time (Table 11).

4.1 Fleet mix. The optimum is 3 ALS and 8 BLS ambulances, serving 82 trips per day (Table 5 and Figure 3). The budget constraint is binding ($₹12,000 + ₹24,000 = ₹36,000$), and one crew member is left spare (25 of 26). The shadow price of the budget is 2.67 trips per extra ₹1,000 per day. The ALS minimum costs 4.67 trips for each additional ALS unit required, which shows the trade-off between serving critical patients and serving more patients in total.

Table 5. Corner points of the feasible region (LPP)

Corner point (x_1, x_2)	Trips per day, Z
(3, 8)	82 (optimal)
(6, 4)	68
(8.67, 0)	52
(3, 0)	18




Figure 3. Graphical solution of the fleet-mix LPP. The shaded area is the feasible region; the optimum is at (3, 8) with $Z = 82$ trips per day.

4.2 Station siting. The set-covering model opens three of the five stations, S2, S4 and S5, and this solution is unique. Every zone is reached within the 8-minute standard; the average time to the nearest open station is 6.25 minutes (Table 6). Stations S1 and S3 can be closed or used as standby posts.

Table 6. Coverage achieved by the optimal stations

Zone	Covering station	Time (min)
Z1	S4	6
Z2	S2	5
Z3	S2	6
Z4	S2	8
Z5	S4	5
Z6	S5	6
Z7	S5	7
Z8	S4	7

Each zone is covered by exactly one open station. If that ambulance is busy, the zone has no backup within 8 minutes. This is the weakness that expected-covering models address (Daskin, 1983).

4.3 Trip allocation. The NWC rule gives a total of 404 response-minutes per day. VAM gives 377 (Table 7), 27 minutes or 6.7% less. The MODI test ($u_1 = 0$, $u_2 = 2$, $u_3 = 1$; $v_1 = 6$, $v_2 = 3$, $v_3 = 5$, $v_4 = 7$) gives positive opportunity costs for all six empty cells (5, 6, 6, 3, 3 and 6), so the VAM solution is optimal and unique. The average response time per trip falls from 6.73 to 6.28 minutes.

Table 7. VAM allocation steps

Step	Highest penalty	Line	Cell allocated	Trips
1	4	Column D2	S2 → D2	12
2	4	Column D3	S5 → D3	16
3	4	Row S2	S2 → D1	3
4	6	Column D1	S4 → D1	15
5	8	Row S5	S5 → D4	9
6	7	Row S4	S4 → D4	5

Ties in the highest penalty were broken by taking the first line; the MODI test confirms that this choice did not affect the optimum. Total = $15 \times 6 + 5 \times 7 + 3 \times 8 + 12 \times 5 + 16 \times 6 + 9 \times 8 = 377$ minutes.

4.4 Dispatch. Row reduction subtracts each row's minimum (5, 10, 6, 6), and column reduction subtracts 1 from column C1. The zeros can then be covered by only three lines (row A2, columns C3 and C4), fewer than four, so the smallest uncovered value, 2, is subtracted from uncovered cells and added at line intersections (Table 8).

Table 8. Final reduced matrix (Hungarian method)

Ambulance	C1	C2	C3	C4
A1	4	5	11	0
A2	0	0	5	6

Ambulance	C1	C2	C3	C4
A3	0	5	0	0
A4	8	0	0	0

Two optimal assignments exist, each totalling 30 minutes. The one chosen ($A1 \rightarrow C4$, $A2 \rightarrow C2$, $A3 \rightarrow C1$, $A4 \rightarrow C3$) has a longest wait of 10 minutes, against 11 for the alternative, so it is fairer to the worst-served patient. The nearest-free-ambulance rule sends A3 to C1, A4 to C2, A2 to C3 and A1 to C4, totalling 35 minutes with a longest wait of 13 (Table 9).

Table 9. Dispatch rule compared with Hungarian assignment

Method	Total time (min)	Longest wait (min)
Nearest free ambulance	35	13
Hungarian method	30	10

The greedy rule fails because it commits A3 to the first call, though A3 is also one of only two ambulances close to C3. Looking at all calls together saves 5 minutes (14.3%).

4.5 Fastest path. Dijkstra's algorithm gives the path $S \rightarrow B \rightarrow A \rightarrow C \rightarrow E \rightarrow H$, taking 15 minutes (Table 10). Going through B to reach A ($2 + 1 = 3$ minutes) beats the direct road $S \rightarrow A$ (4 minutes). The next best paths, the direct road $S \rightarrow A$ (16 minutes) and the detour through F (16 minutes), are close, which gives the driver ready alternatives if a road is blocked.

Table 10. Final Dijkstra labels from station S

Node	S	B	A	C	D	E	F	H
Shortest time (min)	0	2	3	8	10	12	14	15
Previous node	–	S	B	A	C	C	C	E

4.6 Multi-stop route. The nearest-neighbour heuristic builds a 77.4-minute route. It works well early but ends with a long detour from P5 to P7 and back to the depot. Six 2-opt moves remove the crossings and shorten it to 60.0 minutes (Figure 4). Simulated annealing reaches the same 60.0-minute route within about 6,000 iterations (Figure 5): $D \rightarrow P6 \rightarrow P1 \rightarrow P8 \rightarrow P2 \rightarrow P7 \rightarrow P3 \rightarrow P9 \rightarrow P4 \rightarrow P5 \rightarrow D$. Full enumeration confirms it is the exact optimum.

Route improvement: 77.4 → 60.0 min (22.5% shorter)




Figure 4. Nearest-neighbour route (left) and the improved optimal route (right).




Figure 5. Convergence of simulated annealing. The current route often gets worse early on, which lets the search leave poor routes; the best route reaches the optimum and stays there.

In this example with nine stops 2-opt was sufficient. The metaheuristic did not provide any extra improvement. Its benefit increases with size: for paths with 30 or more stops it is not possible to check all options and 2-opt often gets stuck in a local best, which annealing is meant to avoid (Kirkpatrick et al., 1983).

4.7 Summary.

Table 11. Gains from OR methods over simple rules

Decision	Simple rule	OR method	Result	Gain
Fleet mix	–	LPP (graphical)	3 ALS + 8 BLS, 82 trips/day	Budget fully used
Station siting	Open all 5	Set covering ILP	3 stations, 100% coverage	2 stations saved
Trip allocation	NWC: 404 min	VAM + MODI	377 min	6.7%
Dispatch	Nearest free: 35 min	Hungarian	30 min	14.3%
Routing	Nearest neighbour: 77.4 min	2-opt / SA	60.0 min	22.5%

4.8 Discussion. The models fit together as one planning chain. The LPP sets fleet size, the set-covering model places it, the transportation model plans daily workload, and the assignment, shortest-path and routing models guide real-time and scheduled trips. Every model in this paper can be built in Microsoft Excel Solver, so a small operator or a hospital can apply them without special software.

The largest gains came where simple rules ignore interactions: the greedy dispatch rule looks at one call at a time, and the nearest-neighbour route looks one stop ahead. Optimisation helps most exactly where decisions affect each other.

5. Conclusion

This paper showed that the standard undergraduate OR toolkit can support every stage of EMS planning, from fleet size to the final route, on one connected case. On the illustrative district, the optimised decisions cut response time by 6.7% in trip allocation, 14.3% in dispatch and 22.5% in multi-stop routing, while full 8-minute coverage was reached with three stations instead of five. The methods are simple enough to run in a spreadsheet, which makes them practical for smaller operators and hospitals.

Limitations. The data are hypothetical, and travel times are treated as fixed, though real times change with traffic and time of day. The models assume ambulances are always available, while in practice some are busy. Each zone is covered by only one station, so there is no backup coverage.

Future scope. Further work can apply the framework to real call and GPS data from an Indian city, use time-dependent travel times, add backup coverage through the maximal expected covering model, and test simulated annealing, tabu search or genetic algorithms on larger routing instances where they offer more benefit.

Declarations. *Conflict of interest:* The authors declare no conflict of interest.

Funding: This research received no funding.

Data availability: The illustrative data set appears in full in Tables 1 to 4. Figure 2. The Python code used to verify the results can be obtained from the author upon request.

Use of AI tools: AI tools (Claude, Anthropic) were used to assist with drafting, code and figures. The authors reviewed and verified all content and take full responsibility for it.

References

- Aringhieri, R., Bruni, M. E., Khodaparasti, S., & van Essen, J. T. (2017). Emergency medical services and beyond: Addressing new challenges through a wide literature review. *Computers & Operations Research*, 78, 349–368. <https://doi.org/10.1016/j.cor.2016.09.016>
- Bélangier, V., Ruiz, A., & Soriano, P. (2019). Recent optimization models and trends in location, relocation, and dispatching of emergency medical vehicles. *European Journal of Operational Research*, 272(1), 1–23. <https://doi.org/10.1016/j.ejor.2018.02.055>
- Brotcorne, L., Laporte, G., & Semet, F. (2003). Ambulance location and relocation models. *European Journal of Operational Research*, 147(3), 451–463. [https://doi.org/10.1016/S0377-2217\(02\)00364-8](https://doi.org/10.1016/S0377-2217(02)00364-8)
- Church, R., & ReVelle, C. (1974). The maximal covering location problem. *Papers of the Regional Science Association*, 32(1), 101–118. <https://doi.org/10.1007/BF01942293>
- Clarke, G., & Wright, J. W. (1964). Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 12(4), 568–581. <https://doi.org/10.1287/opre.12.4.568>
- Croes, G. A. (1958). A method for solving traveling-salesman problems. *Operations Research*, 6(6), 791–812. <https://doi.org/10.1287/opre.6.6.791>
- Dantzig, G. B., & Ramser, J. H. (1959). The truck dispatching problem. *Management Science*, 6(1), 80–91. <https://doi.org/10.1287/mnsc.6.1.80>
- Daskin, M. S. (1983). A maximum expected covering location model: Formulation, properties and heuristic solution. *Transportation Science*, 17(1), 48–70. <https://doi.org/10.1287/trsc.17.1.48>
- Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1), 269–271. <https://doi.org/10.1007/BF01386390>
- Dorigo, M., & Gambardella, L. M. (1997). Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, 1(1), 53–66. <https://doi.org/10.1109/4235.585892>
- Gendreau, M., Laporte, G., & Semet, F. (1997). Solving an ambulance location model by tabu search. *Location Science*, 5(2), 75–88. [https://doi.org/10.1016/S0966-8349\(97\)00015-6](https://doi.org/10.1016/S0966-8349(97)00015-6)
- Glover, F. (1986). Future paths for integer programming and links to artificial intelligence. *Computers & Operations Research*, 13(5), 533–549. [https://doi.org/10.1016/0305-0548\(86\)90048-1](https://doi.org/10.1016/0305-0548(86)90048-1)
- Hagberg, A. A., Schult, D. A., & Swart, P. J. (2008). Exploring network structure, dynamics, and function using NetworkX. In G. Varoquaux, T. Vaught, & J. Millman (Eds.), *Proceedings of the 7th Python in Science Conference* (pp. 11–15).
- Hillier, F. S., & Lieberman, G. J. (2015). *Introduction to operations research* (10th ed.). McGraw-Hill Education.
- Hitchcock, F. L. (1941). The distribution of a product from several sources to numerous localities. *Journal of Mathematics and Physics*, 20(1–4), 224–230. <https://doi.org/10.1002/sapm1941201224>
- Holland, J. H. (1975). *Adaptation in natural and artificial systems*. University of Michigan Press.
- Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3), 90–95. <https://doi.org/10.1109/MCSE.2007.55>
- Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220(4598), 671–680. <https://doi.org/10.1126/science.220.4598.671>
- Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1–2), 83–97. <https://doi.org/10.1002/nav.3800020109>



Larsen, M. P., Eisenberg, M. S., Cummins, R. O., & Hallstrom, A. P. (1993). Predicting survival from out-of-hospital cardiac arrest: A graphic model. *Annals of Emergency Medicine*, 22(11), 1652–1658. [https://doi.org/10.1016/S0196-0644\(05\)81302-2](https://doi.org/10.1016/S0196-0644(05)81302-2)

Mitchell, S., O'Sullivan, M., & Dunning, I. (2011). *PuLP: A linear programming toolkit for Python*. Department of Engineering Science, The University of Auckland.

National Fire Protection Association. (2020). *NFPA 1710: Standard for the organization and deployment of fire suppression operations, emergency medical operations, and special operations to the public by career fire departments* (2020 ed.).

Rosenkrantz, D. J., Stearns, R. E., & Lewis, P. M., II. (1977). An analysis of several heuristics for the traveling salesman problem. *SIAM Journal on Computing*, 6(3), 563–581. <https://doi.org/10.1137/0206041>

Taha, H. A. (2017). *Operations research: An introduction* (10th ed.). Pearson.

Toregas, C., Swain, R., ReVelle, C., & Bergman, L. (1971). The location of emergency service facilities. *Operations Research*, 19(6), 1363–1373. <https://doi.org/10.1287/opre.19.6.1363>

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., ... Vázquez-Baeza, Y. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>