

Optimal Assignment of Surgeons to Surgical Cases in an Operation Theatre Using the Hungarian Method: An Illustrative Study Based on Indian Hospital Operating Conditions

Sruthi Balaji, Sukanya Pappu, Vedanth Oberoi, Tanush Seth, Vivaan Bhansali

NMIMS ASMSOC

Corresponding author email: sruthibalaji007@gmail.com



<https://doi.org/10.55041/ijssmt.v2i10.029>

Cite this Article: Balaji, S., Pappu, S., Oberoi, V., Seth, T. & Bhansali, V. (2026). Optimal Assignment of Surgeons to Surgical Cases in an Operation Theatre Using the Hungarian Method: An Illustrative Study Based on Indian Hospital Operating Conditions. *International Journal of Science, Strategic Management and Technology*, 02(10), 1-9. <https://doi.org/10.55041/ijssmt.v2i10.029>

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

In any hospital, operation theatres (OTs) are some of the most demanding units in terms of resources. Yet time and again, Indian audit studies put on record the idle theatre time, late starts and cancellations that are the consequence of over-running schedules. We address this by treating the allocation of surgeons to cases as a linear assignment problem and applying the Hungarian method to find a solution (Kuhn, 1955; Munkres, 1957). Our model is an orthopaedic and general surgery list of six cases for six surgeons, with the estimated procedure time in minutes serving as the cost we seek to minimise. Since Indian hospitals do not make surgeon-level time data available to the public, the cost matrix presented here is a simulation based on operational context from published OT audits in India. The paper walks through the entire process in detail, from row and column reduction and the covering-line test, to the θ adjustment and a final optimality check. The results show the optimal call for 460 surgeon-minutes all told. That figure is 22.9 per cent less than one would expect from a random assignment of 597 minutes, and 10.3 per cent under a fixed roster-order at 513 minutes. This unique solution has been borne out by a linear programming formulation run through Excel Solver, complete with dual values and cost ranging. In the discussion we consider how the approach can be applied to unbalanced lists and what limitations are inherent in a deterministic one-to-one model.

Keywords: operation theatre scheduling, assignment problem, Hungarian method, surgeon allocation, hospital operations, India

Optimal Assignment of Surgeons to Surgical Cases in an Operation Theatre Using the Hungarian Method

1. Introduction

In any hospital, the operation theatre is a dual entity: it is among the biggest drains on the budget and at the same time the principal revenue generator. The reason is the confluence of expensive infrastructure and consumables with the need for a highly skilled multidisciplinary team (Pakhare et al., 2022). Consequently, making efficient use of theatre time is a matter of planning that has to be settled long before a patient is brought in. A fundamental part of that planning is determining which surgeon will take on which case; since cases vary in their complexity and surgeons in their experience and subspecialty training, the duration of a procedure is dictated by the one performing it.

There is ample reason to give this kind of planning due diligence, as audits of Indian tertiary care centres would attest. One finds raw theatre utilisation at around 71% (Pakhare et al., 2022), or an average idle time of some 17% even when every planned case is put through (Tyagi et al., 2022). Cancellations owing to a lack of theatre time are also reported (Talati et al., 2015; Vinukondaiah et al., 2000). These two things can happen in tandem because the list is frequently put together without regard for matching the case length to the person doing the work.

We approach the problem of assigning surgeons to cases in this paper as a classical assignment problem, applying the Hungarian method to find a solution. There are three reasons to make the assignment problem the model of choice here.

For one, it is exactly solvable in polynomial time. It also calls for nothing more than a cost matrix that can be put together from a hospital's own records. And the process is sufficiently transparent for a theatre manager to put under review.

The paper offers four distinct contributions: a literature review on the connection between operating room scheduling and assignment models; a fully verified numerical example worked out in full, complete with the covering-line test and θ adjustment; a verification of the outcome via linear programming in Excel Solver and sensitivity analysis; and a benchmarking of the optimal assignment versus two non-optimised options. After the literature review in Section 2, the problem is formulated in Section 3 and the method in Section 4. The data and its limitations are set out in Section 5, followed by the solution in Section 6 and a discussion of results in Section 7. We conclude with the limitations in Sections 8 and 9.

2. Literature Review

2.1 Operating Room Planning and Scheduling

The subject of operating room planning and scheduling has occupied the attention of operational researchers for decades. In a 2010 review, Cardoen and colleagues put the field under the microscope, sorting contributions by such criteria as the patient population in question, the performance metrics and solution techniques employed, how uncertainty is handled and the nature of the planning decisions. What their work makes clear is that the literature is structured around a hierarchy of decisions, ranging from the high level of allocating theatre time to various specialties to the nitty-gritty of sequencing cases for the day.

Samudra et al. (2016) returned to the topic some years later to highlight the pitfalls and challenges of the area, not least the trouble one can have in putting mathematical models into practice.

A number of authors have turned to machine-scheduling analogies when it comes to modelling surgical schedules. Pham and Klinkert (2008), for instance, cast the problem of scheduling surgical cases as a generalised job shop scenario with patients traversing several resources, while Gartner and Kolisch (2014) took on the matter of elective patient flow across the hospital.

More recently in India, Tyagi and his team (2022) found a strong parallel between the parallel machine scheduling of manufacturing and the task at hand; they applied allocation models to fit listed cases into open rooms at a public tertiary eye hospital. Taken together, these papers confirm that machine-scheduling analogies are the norm for modelling theatre decisions, with the assignment problem being regarded as the most straightforward example.

2.2 The Assignment Problem and the Hungarian Method

The assignment problem is one of matching n agents to an equal number of tasks on a one-to-one basis in order to keep the overall cost to a minimum. Kuhn (1955) put forward an efficient way of solving this by making use of concepts found in the work of two Hungarian mathematicians, and he called his approach the Hungarian method. Later, Munkres (1957) presented algorithms for both the transportation and assignment problems and put their computational behaviour under analysis; the procedure he described, with its row and column reductions, covering lines and uniform adjustment of the uncovered elements, has come to be known as the Kuhn-Munkres algorithm. For a contemporary account of these problems and their variants, one can turn to the standard treatment offered by Burkard et al. (2012).

2.3 Unbalanced Assignment Problems

One will not often find the number of cases in a hospital to be the same as the number of surgeons. Such an assignment problem, where agents and tasks are unequal in number, is termed unbalanced. The conventional approach found in textbooks is to make the matrix square by introducing dummy rows or columns at zero cost; this has the effect of matching some tasks to a non-existent agent, leaving them unexecuted for that round (Kumar, 2006). Kumar (2006) put forward an alternative: a modified method to break the unbalanced problem down into balanced sub-problems so every job gets done. A like-minded approach was taken by Amaliah et al. (2013) in assigning surgeons to disaster victims, where they partitioned an 8×5 cost matrix into two balanced ones.

These are pertinent to the present discussion since a theatre list will normally have more cases than surgeons. One can take the practical view that a dummy-based solution is workable in successive rounds, with one case apiece being assigned to each surgeon until the docket is cleared.

2.4 Empirical Evidence from Indian Operation Theatres

There is a useful if limited audit literature on OT functioning in India. In a prospective study of general surgery at a teaching hospital, Vinukondaiah et al. (2000) noted that while the theatre was active 279 days of the year and 91.5% of operating time was put to use, 14.9% of cases were called off. The lack of time was the prime reason for these cancellations, accounting for 65.2% of them. An examination of a major tertiary institute in North India pointed to restrictions on general anaesthesia time as the chief cause of cancellation (Talati et al., 2015).

More recent figures come from an audit at ESIC Medical College and Hospital in Hyderabad. Pakhare et al. (2022) recorded a raw utilisation of 71.31% of resource time and an average turnover of 13.73 minutes. Of 126 elective lists, 23.01% saw an overrun of 35.95 minutes on average, and 8.99% were cancelled. At a public-sector ophthalmology hospital, Tyagi et al. (2022) observed that even when all planned cases were accommodated, the theatre was idle for an average of 17% of the day, and suggested that scheduling based on procedure time would stem over-runs and the attendant cancellations.

Taken together, the audits suggest Indian OTs are beset by both idle time and time-related cancellations. Both point to lists that do not accord with realistic case durations, making a strong case for straightforward optimisation tools.

2.5 Research Gap

What has been surveyed is either a broad review of scheduling models or an account of OT performance; there is little in the way of a fully worked, reproducible assignment example set in an Indian context. Even the published work of Amaliah et al. (2013) on the application of a modified Hungarian method to disaster management omits any detailed exposition of the covering-line and θ steps. This paper is intended to fill that void with a complete worked example relevant to Indian operating conditions.

3. Problem Formulation

Take the example of a morning elective with n surgeons and an equal number of cases to be run in parallel theatres, one per surgeon. We denote by c_{ij} the time in minutes surgeon i is estimated to require for case j , and x_{ij} as 1 should that surgeon be assigned to it (0 if not). The model then becomes:

Minimise $Z = \sum_i \sum_j c_{ij} x_{ij}$

subject to $\sum_j x_{ij} = 1$ for all i ; $\sum_i x_{ij} = 1$ for all j ; $x_{ij} \in \{0, 1\}$.

By these constraints, each case is the work of exactly one surgeon and each surgeon operates on one case only. In this way we are minimising total time, which serves to mitigate the risk of overrun since cancellations of that nature are a function of total surgeon-time (Pakhare et al., 2022; Vinukondaiah et al., 2000).

4. Methodology: The Hungarian Method

We apply the Kuhn–Munkres procedure to the square cost matrix in its standard form (Burkard et al., 2012; Kuhn, 1955; Munkres, 1957) through the following steps:

Step 1. Row reduction. From every element in a row, subtract the row's smallest value.

Step 2. Column reduction. Do the same for each column. At this point there will be at least one zero in every row and column.

Step 3. Covering-line test. Use the fewest horizontal and vertical lines possible to cover all zeros. Should the count of lines be n , an optimal assignment can be found among the zeros and one may proceed to Step 5.

Step 4. θ adjustment. Identify the smallest uncovered element as θ . This is subtracted from all uncovered elements and added to any at an intersection of two lines; those under a single line are left alone. Then return to Step 3.

Step 5. Assignment. Pick n independent zeros, one per row and column, beginning with the rows or columns having but a single zero.

Step 6. Optimality check. Add up the original costs for the chosen positions. If the sum matches the dual objective (the aggregate of all row and column reductions plus θ times the increase in lines required at each adjustment), optimality is confirmed.

The combined reductions from the first two steps provide a lower bound on the optimal cost, as any full assignment draws one element from each row and column. Each θ adjustment lifts the bound by $\theta \times (n - k)$ where k is the number of covering lines. Once $k = n$ the bound is met by the selected zeros.

5. Data and Case Setting

5.1 Data Availability

It is institutional policy and a matter of patient confidentiality that Indian hospitals do not make their surgeon-level procedure-time and cost matrices public. For this reason the present study makes use of a simulated dataset for illustration. The authors have put together plausible figures for the cost matrix to show how a specialist will tend to be quicker with a case in his own subspecialty than with one outside it. These are not to be taken as measured performance data from any particular hospital or person. While the Indian OT audits in Section 2.4 are referenced to give some operational context in terms of turnover, utilisation and so forth, they do not inform the matrix entries.

5.2 Cases and Surgeons

Table 1 gives the six elective cases from a joint general surgery and orthopaedic session. With six surgeons of varying subspecialty backgrounds on hand, the aim is to keep the overall surgical time to a minimum across the six theatres. Cost here is simply the estimated minutes from incision to closure.

Table 1

Surgical Cases and Surgeons in the Illustrative Session

Case	Procedure	Surgeon	Profile
C1	Laparoscopic cholecystectomy	S1	Laparoscopic general surgeon
C2	Open inguinal hernioplasty (mesh)	S2	General surgeon
C3	Total knee arthroplasty	S3	Arthroplasty specialist
C4	ORIF of long-bone fracture	S4	Orthopaedic trauma surgeon
C5	Laparoscopic appendectomy	S5	Sports and arthroscopy surgeon
C6	Arthroscopic ACL reconstruction	S6	Senior generalist

Note. All entries are illustrative. ORIF = open reduction and internal fixation; ACL = anterior cruciate ligament.

Table 2

Illustrative Cost Matrix: Estimated Surgical Time (Minutes)

	C1	C2	C3	C4	C5	C6
S1	55	48	150	140	42	135
S2	62	50	145	135	50	130
S3	105	95	105	118	90	115
S4	100	90	125	100	88	120
S5	98	92	130	125	85	95
S6	70	58	135	128	55	118

Note. Rows are surgeons and columns are cases. Values are simulated for demonstration.

6. Solution

6.1 Step 1: Row Reduction

The row minima are 42, 50, 90, 88, 85, and 55 for S1 to S6, respectively. Subtracting each minimum from its row gives Table 3. The sum of the row reductions is 410.

Table 3

Matrix After Row Reduction

	C1	C2	C3	C4	C5	C6
S1	13	6	108	98	0	93
S2	12	0	95	85	0	80
S3	15	5	15	28	0	25
S4	12	2	37	12	0	32
S5	13	7	45	40	0	10
S6	15	3	80	73	0	63

Note. Shaded cells are zeros.

6.2 Step 2: Column Reduction

The column minima of Table 3 are 12, 0, 15, 12, 0, and 10 for C1 to C6. Subtracting them gives Table 4. The sum of the column reductions is 49, so the lower bound so far is $410 + 49 = 459$ minutes.

Table 4

Matrix After Column Reduction

	C1	C2	C3	C4	C5	C6
S1	1	6	93	86	0	83
S2	0	0	80	73	0	70
S3	3	5	0	16	0	15
S4	0	2	22	0	0	22
S5	1	7	30	28	0	0
S6	3	3	65	61	0	53

Note. Shaded cells are zeros.

6.3 Step 3: The Covering-Line Test

Five lines are sufficient to cover the zeros in Table 4, namely column C5 and the rows for S2 through S5. Since a full assignment would require six, we have not yet reached an optimal matrix. This is apparent from the fact that surgeons S1 and S6 have a zero exclusively in column C5, precluding either from being assigned there. In all, there are but five independent zeros.

6.4 Step 4: θ Adjustment

We turn to the uncovered elements, which are found in columns C1, C2, C3, C4 and C6 of rows S1 and S6. The minimum among them is $\theta = 1$ at (S1, C1). By subtracting 1 from each of these and putting 1 into the four elements where the lines cross (C5 with S2, S3, S4 and S5), we arrive at Table 5. As a result the lower bound is increased by $\theta \times (6 - 5)$ or 1, to 460 minutes.

Table 5

Matrix After the θ Adjustment ($\theta = 1$)

	C1	C2	C3	C4	C5	C6
S1	0	5	92	85	0	82
S2	0	0	80	73	1	70

	C1	C2	C3	C4	C5	C6
S3	3	5	0	16	1	15
S4	0	2	22	0	1	22
S5	1	7	30	28	1	0
S6	2	2	64	60	0	52

Note. Shaded cells are zeros.

6.5 Step 5: Assignment

With six lines required to cover the zeros, we have an optimal assignment on our hands. Three of them are a given: S6 with its lone zero at C5, S5 at C6 and S3 at C3. That leaves S1 to be put in column C1, as there would be no zero for it otherwise. As for the rest, S4 is made to match with C4 and S2 with C2. In the end the assignment is unique.

Table 6

Optimal Assignment and Time

Surgeon	Case	Case ID	Time (min)	Share of total (%)
S1	Laparoscopic cholecystectomy	C1	55	12.0
S2	Open inguinal hernioplasty	C2	50	10.9
S3	Total knee arthroplasty	C3	105	22.8
S4	ORIF of long-bone fracture	C4	100	21.7
S5	Arthroscopic ACL reconstruction	C6	95	20.7
S6	Laparoscopic appendectomy	C5	55	12.0
Total			460	100.0

6.6 Step 6: Optimality Check

With the selected cells incurring original costs of 55, 50, 105, 100, 95 and 55, the total comes to 460 minutes. This figure is in line with the lower bound derived from the row (410) and column (49) reductions along with the θ adjustment of 1, indicating an optimal assignment. A standard linear sum assignment routine was run on the matrix as a matter of course; it produced the identical 460 minute total and the same assignment.

6.7 Verification Using Linear Programming (Excel Solver)

Being a particular type of linear program, the assignment problem can be verified by means of a general-purpose LP solver. The model set out in Section 3 was put through Microsoft Excel's Solver add-in (Fylstra et al., 1998), which employed the Simplex LP method to find a solution.

In this setup, the $36 x_{ij}$ decision variables are contained in a 6×6 block of changing cells, while the objective function is the SUMPRODUCT of the cost matrix and that block. There are twelve equality constraints to see that each row and column sum to one, and the variables are non-negative. No need was made for integrality constraints since the constraint

matrix of an assignment problem is totally unimodular, making the LP optimum integral by nature (Burkard et al., 2012).

As Table 7 shows in its Answer Report format, the Solver output is consistent with what was found via the Hungarian method. All twelve constraints were binding and the six surgeon–case pairs have $x = 1$, for a combined time of 460 minutes. The two approaches thus concur on the optimal value and the assignment itself.

Table 7

Solver Answer Report Summary

Item	Cell	Original value	Final value
Objective: total time (min)	\$B\$22	0	460
$x(S1, C1)$	\$B\$13	0	1
$x(S2, C2)$	\$C\$14	0	1
$x(S3, C3)$	\$D\$15	0	1
$x(S4, C4)$	\$E\$16	0	1
$x(S5, C6)$	\$G\$17	0	1
$x(S6, C5)$	\$F\$18	0	1
Surgeon constraints (6): row sum = 1	\$H\$13:\$H\$18		All binding
Case constraints (6): column sum = 1	\$B\$19:\$G\$19		All binding

Note. All other decision variables have a final value of 0.

A second level of verification is provided by the sensitivity analysis. Table 8 sets out the dual values, or shadow prices, for the twelve constraints. In keeping with LP duality, these total 460 to match the optimal objective value. One will find they are identical to the cumulative row and column reductions from the Hungarian method (to wit, the row reductions with an added 1 for the θ adjustment, and the column reductions minus 1 in the case of C5). As for the reduced costs, subtracting the dual values of the relevant row and column from the original figures yields the same results as in Table 5.

Table 8

Dual Values (Shadow Prices) of the Assignment Constraints

Surgeon constraint	Dual value	Case constraint	Dual value
S1	43	C1	12
S2	50	C2	0
S3	90	C3	15
S4	88	C4	12
S5	85	C5	–1
S6	56	C6	10
Sum	412	Sum	48

Note. Total of all dual values = $412 + 48 = 460$ minutes. This LP is primal-degenerate, so the dual values are not unique; another optimal dual solution, such as the one printed in a Solver Sensitivity Report, may differ cell by cell but must sum to 460.

The figures in Table 9 show the extent to which any given time estimate can be increased without altering the optimal assignment. These are exact ranges, determined by re-solving the problem with the pair in question left out. As for a

reduction in the time of a selected pair, that will not affect the assignment in any way; hence there is no limit to the allowable decrease.

Table 9

Sensitivity of the Selected Assignments to Their Time Estimates

Selected pair	Time (min)	Allowable increase (min)	Allowable decrease (min)
S1 – C1	55	2	No limit
S2 – C2	50	2	No limit
S3 – C3	105	38	No limit
S4 – C4	100	38	No limit
S5 – C6	95	45	No limit
S6 – C5	55	2	No limit

Note. Allowable increase is the largest rise in the pair's time for which the optimal assignment stays the same.

The ranges show that the solution is not equally robust everywhere. The three orthopaedic assignments (S3–C3, S4–C4, S5–C6) tolerate increases of 38 to 45 minutes. The three assignments among S1, S2, and S6 tolerate only 2 minutes, because these surgeons have close substitutes for one another on the general surgery cases. In practice, errors of a few minutes in estimated times could therefore reverse those three pairings, and a theatre manager might treat them as interchangeable. Conversely, the smallest time reduction that would bring any non-selected pair into the optimal solution is 2 minutes.

7. Results and Discussion

To judge the value of the optimal assignment, two benchmarks were computed. The first is a fixed roster-order assignment in which surgeon i operates on case i (S1 to C1 through S6 to C6). Its cost is $55 + 50 + 105 + 100 + 85 + 118 = 513$ minutes. The second is the expected cost of a uniformly random assignment, which equals the sum of all matrix entries divided by n ($3,582 \div 6 = 597$ minutes). Table 10 compares the three.

Table 10

Comparison of Assignment Strategies

Strategy	Total time (min)	Mean per case (min)	Saving of Hungarian (%)
Random assignment (expected)	597	99.5	22.9
Fixed roster order (Si to Ci)	513	85.5	10.3
Hungarian method (optimal)	460	76.7	—

Note. Savings are the reduction in total time achieved by the Hungarian method relative to each strategy.

In the course of a single morning session, the optimal assignment puts in 53 minutes' time compared with the roster order and 137 minutes over a random one. There is a clinical logic to it: the arthroplasty specialist is on the knee replacement, the trauma surgeon on the fracture and the sports surgeon on the arthroscopic ligament work. What is less expected is the senior generalist (S6) being given the appendectomy in preference to an orthopaedic case. The reason is that S6 has a marked disadvantage on the orthopaedic side and the appendectomy is the least costly option left; such non-obvious trade-offs are where a systematic approach has the edge over informal allocation.

One should view these savings in their operational context. Indian audits have put average idle time at 17 per cent of theatre time (Tyagi et al., 2022) and overruns at some 36 minutes for lists that go over (Pakhare et al., 2022). A 53-minute improvement on the roster is therefore on a par with a typical over-run, which would seem to make a case for

better matching as a way to cut time-related cancellations. This is more of a hypothesis to be tested empirically than a firm conclusion, given the simulated nature of the matrix.

7.1 Extension to Unbalanced Lists

Teaching hospitals often have more cases on the list than surgeons. To deal with this, the matrix can be padded with zero-cost dummy surgeons and the Hungarian method run in rounds until every case is placed (Burkard et al., 2012; Kumar, 2006), or it can be broken down into balanced sub-matrices (Kumar, 2006; Amaliah et al., 2013). Since each round is a theatre slot in the day, they can be sequenced to account for the 11 to 17 minutes of turnover between cases found in the Indian data (Pakhare et al., 2022).

8. Limitations and Future Research

The findings are subject to a number of caveats. The cost matrix is for illustration only and the results should not be taken as evidence for any particular hospital. The model is also deterministic, unlike the uncertain and skewed durations of actual procedures. It is a one-to-one assignment that does not factor in anaesthetists, nurses or room availability, all of which are central to the integrated models in the literature (Cardoen et al., 2010; Samudra et al., 2016). Nor does it allow for emergencies or surgeon preference, and minimising total time is no guarantee of fair workloads or fewest cancellations.

There are four avenues for future research. These include working with a hospital to derive real time matrices from OT registers; using robust or stochastic extensions to cover duration variability; adopting generalised assignment or job-shop formulations for multi-case lists along the lines of Pham and Klinkert (2008); and running a pilot with before-and-after audits to see if the savings do in fact lead to fewer cancellations (Pakhare et al., 2022).

9. Conclusion

We have demonstrated how the operation theatre's surgeon-to-case decision can be treated as an assignment problem and solved exactly by means of the Hungarian method. A six-by-six example was worked through in full to verify optimality via the dual bound and an Excel Solver linear programming formulation, covering the θ adjustment and the covering-line test. The resulting 460-minute solution was unique and some 22.9 per cent below a random allocation and 10.3 per cent under a fixed roster. With the idle time and late starts recorded in Indian OT audits, there is merit in putting this kind of straightforward optimisation to the test with real data. While the one-to-one structure is a limitation, it offers a sound basis for the development of more complex scheduling models.

References

- Amaliah, B., Widyanto, M. R., Purwanto, D., Armand, F., Iqbal, M., & Verdianto, S. (2013). A modified Hungarian method for unbalanced assignment doctor problem in disaster management. In Y. Agusta (Ed.), *Proceedings of the 3rd International Workshop on Soft Computing and Disaster Control (SocDic 2013)*. STMIK STIKOM Bali.
- Burkard, R., Dell'Amico, M., & Martello, S. (2012). *Assignment problems* (Revised reprint). Society for Industrial and Applied Mathematics.
- Cardoen, B., Demeulemeester, E., & Beliën, J. (2010). Operating room planning and scheduling: A literature review. *European Journal of Operational Research*, 201(3), 921–932. <https://doi.org/10.1016/j.ejor.2009.04.011>
- Fylstra, D., Lasdon, L., Watson, J., & Waren, A. (1998). Design and use of the Microsoft Excel Solver. *Interfaces*, 28(5), 29–55. <https://doi.org/10.1287/inte.28.5.29>
- Gartner, D., & Kolisch, R. (2014). Scheduling the hospital-wide flow of elective patients. *European Journal of Operational Research*, 233(3), 689–699.
- 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>
- Kumar, A. (2006). A modified method for solving the unbalanced assignment problems. *Applied Mathematics and Computation*, 176(1), 76–82. <https://doi.org/10.1016/j.amc.2005.09.056>



- Munkres, J. (1957). Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1), 32–38. <https://doi.org/10.1137/0105003>
- Pakhare, V., Gopinath, R., Surya Dhanalakshmi, S. K., Nanda, A., Kanojia, N., & Venu, P. (2022). Audit of operation theater time utilization with perspective to optimize turnaround times and theater output. *Journal of Anaesthesiology Clinical Pharmacology*, 38(3), 399–404. https://doi.org/10.4103/joacp.joacp_398_20
- Pham, D.-N., & Klinkert, A. (2008). Surgical case scheduling as a generalized job shop scheduling problem. *European Journal of Operational Research*, 185(3), 1011–1025.
- Samudra, M., Van Riet, C., Demeulemeester, E., Cardoen, B., Vansteenkiste, N., & Rademakers, F. E. (2016). Scheduling operating rooms: Achievements, challenges and pitfalls. *Journal of Scheduling*, 19(5), 493–525.
- Talati, S., Gupta, A. K., Kumar, A., Malhotra, S. K., & Jain, A. (2015). An analysis of time utilization and cancellations of scheduled cases in the main operation theater complex of a tertiary care teaching institute of North India. *Journal of Postgraduate Medicine*, 61(1), 3–8.
- Tyagi, M., Tyagi, P. K., Singh, S., Satpathy, S., Kant, S., Gupta, S. K., & Singh, R. (2022). Allocation scheduling leads to optimum utilization of operation theater time. *Medical Journal Armed Forces India*, 78(Suppl. 1), S163–S171. <https://doi.org/10.1016/j.mjafi.2020.09.005>
- Vinukondaiah, K., Ananthakrishnan, N., & Ravishankar, M. (2000). Audit of operation theatre utilization in general surgery. *National Medical Journal of India*, 13(3), 118–121.