
Dynamic Stakeholder-Centred Surgical Scheduling



Hospitals must keep operating theatres productive while protecting staff wellbeing and delivering timely, effective care. That balance is difficult when schedules are disrupted by variable procedure times, cancellations and limited specialist availability. A scalable methodology sets out how to optimise surgical timetables dynamically across competing interests. It combines a stepwise framework with probabilistic modelling to handle uncertainty and align decisions with the priorities of administrators, clinicians, students and patients. A case application in a specialist hospital setting demonstrates how the approach reduces delays, improves workload balance and strengthens operational resilience without sacrificing quality or safety.

From Fragmented Goals to a Unified Objective

Many optimisation efforts in healthcare scheduling concentrate on a narrow slice of the problem, targeting either cost, utilisation or staff workload in static settings. The framework described here addresses the gap by unifying objectives across stakeholders and updating decisions as conditions change. It integrates Bayesian probability modelling with a utilitarian optimisation function that aggregates net benefits for owners, users and clients, expressed in comparable monetary terms that encompass economic and social costs. Rather than privileging one metric, it weighs income, operating and maintenance expenses, depreciation, training, administrative overheads and salaries against impacts such as waiting time, accidents, stress and illness. By treating these as components of a single objective, the method balances throughput with staff conditions and patient value, then recalculates as new information arrives.

The 12-step process begins with scope definition, resources and service types, advances through stakeholder impact hierarchies and objective formulation, and culminates in generating candidate schedules, evaluating net benefit and iterating when service requests change. It is designed to remain responsive to unexpected events so scheduling aligns continuously with operational reality.

Modelling Uncertainty Where It Matters

Uncertainty is modelled explicitly for the variables that most influence outcomes: probability of surgical success, expected waiting time from shift delays, accident probability per person, probability of illness and the likelihood of excessive stress. Each variable is linked to contextual conditions and timeframes, then estimated using historical data and expert input via Bayesian networks that support posterior updates when new data arrive. This embeds learning into the schedule rather than treating variability as noise.

Must Read: [Optimising Acute Care Efficiency with AI-Driven Patient Flow Optimisation](#)

Surgical success depends on the procedure type, surgeon experience, daily workload and room suitability. Records from 2023 indicate an average success rate of 81%, with adjustments for context. High daily workload reduces success by 2% while low workload increases it by 2%. Advanced-level surgeons add a 1% uplift whereas early-career surgeons apply a 1% reduction. Room choice also matters. A versatile room supports better outcomes, while spaces designed for specific case types can disadvantage others, and smaller rooms can slightly disadvantage larger or more complex cases. These conditional relationships are encoded so that schedule decisions reflect how staffing, rooming and case mix interact in practice.

Operational realities are folded into the model. Turnover times and overrun risk are considered, including a 10% chance of a 0.5-hour overrun in complex procedures. A 5% same-day cancellation probability prompts re-optimisation, and shift caps at eight hours incorporate overtime penalties. These constraints anchor the optimisation in feasible theatre use and sustainable staffing patterns.

Proof of Concept in a Complex Surgical Day

The case setting is a specialist hospital service that must combine flexibility for emergencies with efficient routine care under limited specialist capacity. The demonstration day comprises five planned procedures spanning different techniques. The service catalogue covers nine surgical categories with indicative durations. Resources include five rooms with distinct suitability, six surgeons and five nurses, with surgeons grouped by experience because not all can perform every procedure. For example, only two surgeons are qualified for ophthalmic procedures, and some early-career surgeons are not assigned to specific complex work.

Room characteristics shape feasible allocations. Some rooms are equipped for larger or more complex cases, one room is dedicated to a specific patient type due to specialised equipment, and smaller rooms are optimised for routine procedures. A central room is the most versatile and can host the full range of case types. Candidate schedules map these constraints, linking each case with eligible rooms, surgeons and nurses while respecting capacity limits and experience requirements.

The optimisation evaluates multiple candidate timetables against the composite objective, selecting the schedule that maximises expected net benefit given contextual probabilities. Performance is then tracked against actuals, with the procedure repeated when conditions change. In the worked example, a new abdominal case arrives late in the morning, triggering a full reschedule under the same framework. This iterative loop demonstrates how the method absorbs new demand without losing alignment with stakeholder goals.

Results benchmarked against traditional scheduling show improvements in total net benefit on average by 7.80%, rising to 15.28% in the worst-case baseline scenario. In the initial run, one candidate timetable delivered the best balance across owners, users and clients. After the first update, a revised variant maintained the highest total net benefit while ensuring completion in a single day, which avoided overtime and associated costs. These gains reflect fewer delays, better workload distribution and prudent use of rooms and skills.

A dynamic, stakeholder-centred approach to surgical scheduling can improve the net benefit of hospital operations by modelling uncertainty where it matters, unifying competing objectives and iterating as conditions change. By tying success probabilities, delays, accidents, stress and illness to the real context of rooms, workload and experience, and by embedding practical constraints such as overruns, cancellations and shift limits, schedules become both feasible and resilient. The case application shows measurable gains over traditional methods, with average and peak improvements in net benefit, and indicates a route to scale the methodology across services to support decision-making that serves patients, staff and administrators together.

Source: [Healthcare Analytics](#)

Image Credit: [iStock](#)

Published on : Wed, 10 Sep 2025