Rapidly optimising healthcare resourcing schedules in dynamic situations

L.Law and J.Shimell. Autonomy & Spatial Modelling Team – ***tpgroup***.

Version 1.0, 12th May 2020

Abstract

One of the key elements of the UK Government's strategy to fight the COVID-19 outbreak was reducing the potential for such a pandemic virus to overwhelm any part of the National Health Service (NHS). Such an impact would have been both instant and long lasting in the exponential increase in the number of deaths caused by COVID-19 and others conditions, diseases or health conditions that could not be treated as a result.

This white paper describes how the Optimiser tool can rapidly and repeatedly provide suggested optimal schedules for limited healthcare resources to keep up with a dynamic, complex COVID-19 (or similar) outbreak, across multiple hospitals or health settings. The tool could support the planning and logistical expertise of healthcare planners in optimising the use of the staff, equipment and consumables of an already constrained healthcare system.

The output of such work would seek to address future similar scenarios. However, if the relevant data and support were available, the project could be delivered within the next six months under the current and normal working restrictions to aid the recovery phase of COVID-19 where hospitals will need to maintain a more normal general workload - while maintaining atypical measures in dealing with COVID-19 outbreaks.

The need

The National Health Service (NHS) has a legal obligation, as defined in the Civil Contingencies Act 2004 and NHS Act 2006, to plan and prepare in order to effectively respond to a range of specified and unspecified emergencies that could affect health and patient care. Delivery of this role is through multi-disciplinary and multi-institutional Emergency Preparedness, Resilience and Response[1] (EPRR) teams. The current COVID-19 outbreak has seen many of these plans enacted, including the enactment of military aid to the civil authorities, and local resilience forums (LRF's). These are multi-agency partnerships made up from local public services including the emergency services, local authorities, the NHS, the Environment Agency and others helping coordinate the responses efforts at their local level.

To generate a plan to provide medical support, consideration needs to be given to grades of available staff, their skills, an appropriate shift pattern, equipment availability, equipment or resource lead times, bed availability fluctuations, transport times and availability of ambulances and associated staff for that transport etc. All of this needs to be considered, and a plan created that will crucially be constantly updated as the situation changes dynamically. This presents an enormous challenge for health planners.

Healthcare resource optimisation is a large-scale problem with many involved factors that are spread both temporally (seasonal fluctuations and demands, along with routine procedures and care programmes), and geographically (staff and equipment movements). The NHS is the largest employer in the UK, heavily reliant on professional clinical staff that account for half of the workforce. In 'normal' circumstances the organisation works with an overall shortfall of staff. Approximately 1 in 12 posts are estimated to be vacant in hospital and

---

[1] The Government's Emergency Preparedness publication provides guidance on Part 1 of the Civil Contingencies Act 2004, its associated Regulations and non-statutory arrangements; focussing on the pre-emergency (planning) phase. Emergency Response and Recovery describes the multi-agency framework for responding to, and recovering from, emergencies in the UK.

community services, with the largest area of vacancies being nursing and midwifery related. This only increases the complexity associated with achieving appropriate staffing levels in a normal operational scenario, without the added complication of dynamically changing constraints due to staff getting sick and/or having to shield as they are vulnerable – which is being experienced in the current COVID-19 outbreak.

As the COVID-19 outbreak moves from the response phase to the recovery phase, hospitals and clinical commissioning groups will face a hard decision on the balance of resources required for COVID-19 wards against the more normal operational requirements of the healthcare system.

While there are many modelling and simulation tools currently available to planners, and patient pathways are well understood, these can only **predict the outcome** of a given plan of use for limited resources. There are currently no known optimisation solutions in the healthcare space that can make **decisions on how to optimise limited resources**. In other words, existing systems can tell health planners the likely outcome of a plan, they cannot tell them what plan is best.

Our solution

tpgroup recognised that traditional management processes and systems lack the flexibility and responsiveness to deal with complex and dynamic environments with multiple (often autonomous) resources - thus ***tpgroup's*** Optimiser tool was developed. Optimiser is a dynamic and rapid decision support tool for organisations that need to plan capacity and optimise the availability of critical resources within their enterprise.

Optimiser was initially developed to optimise operation and maintenance in the offshore wind industry where it sought to achieve optimal schedules for a wide variety of assets and resources in multiple locations, in **minimal time**, for a **dynamically changing environment**. This was whilst balancing a number of factors – cost, energy yield, customer KPIs etc. Following this initial work, ***tpgroup*** are now working with industry leading data experts in the onshore oil & gas industry to rework Optimiser to be deployed for use within the oil and gas sector in their move to become a smarter industry.

Drawing on our other use cases, Optimiser could be adapted and developed into an innovative healthcare solution to support health planners, taking dynamic short-term demand forecasts for the allocation of staff and resources as an input and to generate and output optimal staff and resource schedules based on multiple chosen constraints, in minimal time**. The use of Optimiser would enable a planning team to balance the demand and resources across the care centres available (including NHS hospitals and field hospitals) to reduce the likelihood of any one care facility being overwhelmed, thus improving patient outcomes, saving lives and reducing stress on all staff – both planning and operational.**

Use Case

The allocation of resources across multiple centres of care delivery within a clinical commissioning group (CCG) is complex - balancing current and forecast patient demand against the availability of critical and often limited resources.

A healthcare planner needs to consider a wide range of factors: the current state of the system; the availability of multiple grades of staff and their shift patterns; equipment availability and lead times; bed availability; transport times and availability of ambulances and associated staff for that transport etc. All of this needs to be considered and a plan created and, crucially, constantly updated as the situation changes. This presents an enormous challenge for any individual. Figure 1 below sets this out pictorially across a group of hospitals.

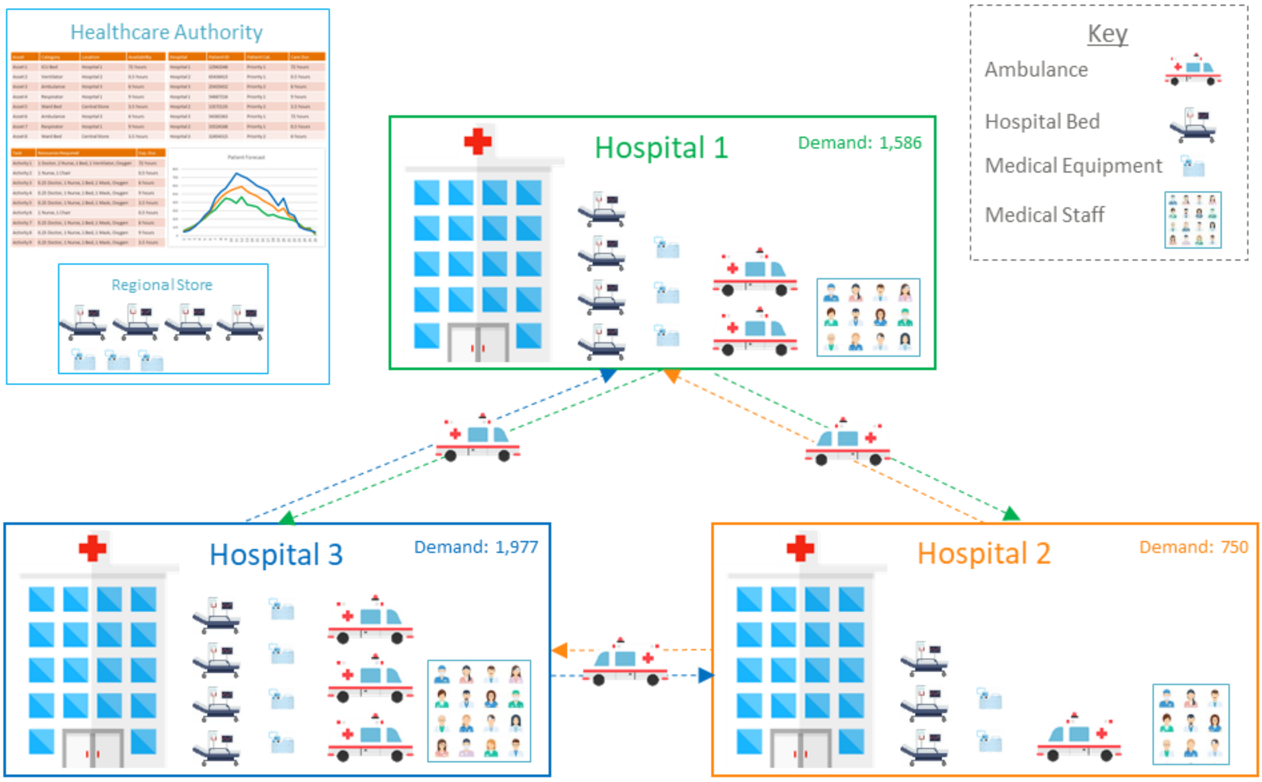

Figure 1 Balancing resources across multiple hospitals in a group

Optimiser would look to balance the allocation of limited resources, such as ventilators and respirators or specialist ICU staff, as during the COVID-19 crisis, across multiple healthcare resources – using modelling information from the NHS and government. It would identify areas of the country where additional beds, and potential temporary field hospital facilities may be required, as well as managing the commissioning and build of such a facility - through to an operational stage and its decommissioning. It is a tool that could be used for a single hospital, and also be used in conjunction with groups of healthcare resources within the wider local NHS CCG's to bring a local joined up approach.

Usable Optimisation – Excellent now rather than perfect too late

Usually generating such optimisations by the traditional Exact[2] methods would require large amounts of computational power, take too much time for the resulting output to be of use, or find that the situation / input data had changed during that time. At the scales likely to be found in the healthcare use case set out here, traditional exact optimisation approaches struggle to keep up with the dynamic nature of the problem. These approaches typically have long run times, meaning such alternative systems cannot respond rapidly enough to unfolding events with no solution is available until the final solution is produced.

As an alternative, heuristic approaches use a semi random, guided exploration of a decision space, pursuing promising solutions to rapidly find near optimal solutions. Whilst a heuristic approach can explore a single dimension (or factor) of optimisation, a metaheuristic approach can explore multiple dimensions simultaneously. Optimiser is an artificial intelligence (AI) based system that uses a metaheuristic approach to explore large-scale, multidimensional decision spaces, such as that presented in this use case, through an iterative master process

---

[2] Optimisation approaches are typically divided into two types, Exact (sometimes called Direct) - finding the optimum (e.g. Lowest Cost) possible solution, by exploring all possible options; or Heuristic (Defined as: proceeding to a solution by trial and error or by rules that are only loosely defined.)

that guides and modifies subordinate heuristics to produce high quality solutions (or decisions), and encouraging them to find global, rather than local, optima. In practice this allows the system to find a balanced optimisation across multiple factors, rather than maximising or minimising a single factor. For example the system could find the best balance of cost to production, rather than a solution that just maximises one of those factors.

In a comparative test using an offshore renewable energy data set, Optimiser was compared to a traditional multi integer linear programming (MILP) exact optimisation approach, producing an operations and maintenance schedule for a set of wind turbines. Optimiser produced a result within 0.01% of that computed by the MILP exact method in less than 1% of the run time – over 500 times faster (15 seconds compared to 3 hours). **In practice, a close to optimal solution with rapid computation is more useful than an exact, but slow solution**.

Optimiser is based on the more rapid meta-heuristic approach, and can be run as frequently as situations change, always using the current state of the system as its starting point. This approach is also "interruptible", as in it finds a solution rapidly and then improves upon it. This allows a runtime to be set, at the end of which it will produce an answer. The longer the system runs, the closer to optimal it will be. This "interuptability" and rapid runtime allows several benefits compared to other approaches:

- The system can be rerun to account for changing conditions
- The system can run multiple times to explore different scenarios
- The user can set a runtime based on their current needs, leaving it running longer if they have time, or requesting a shorter run if there is urgency.

Data Requirements

The typical key data items that Optimiser would consider in this use case are:

- Availability of assets per centre and in central stores. This will include a forecast availability for the period of interest – for example when a delivery of respirators is expected:
  - Staff of different skills/grades – incl. shift patterns
  - Beds of different types (e.g. ICU)
  - Respirators
  - Transport (Ambulances)
- Locations of centres and the duration of travel between them
- Forecast of patient numbers per hospital/region along with their:
  - Category of care required (e.g. ICU)
  - Assumed duration of care
  - Likelihood of moving category (i.e. moving to ICU)
  - Survival chance
- Activities required and the associated activities and resources required:
  - Patient care (by category e.g. ICU)
  - Transport – patient
  - Transport – staff
  - Transport – equipment

There would be no need to use anything other than anonymised data. Some of this anonymised data could be taken from historic record - some will be calculated, and some will be assumed. This latter element presents a potential problem for all optimisation and modelling processes, as if these assumptions are incorrect, then the optimisation will be flawed. The old data science adage of "Rubbish in, Rubbish out" is as true in today's big data world as it has ever been. ***tpgroup*** are working on machine learning approaches that could potentially help to adjust these assumptions in real-time by monitoring the reality of the situation on the ground, thus improving the accuracy of the assumptions used, which could be a part of future exploitation of the system.

Output

The output of Optimiser would be an **optimised schedule for the use of key assets for the period of interest** to maximise (or minimise) for a chosen factor – for example to maximise the number of patients cared for, or minimise the loss of life. The primary aim is to prevent any one hospital facility being overwhelmed or ambulances being turned away from facilities with critical patients as has happened elsewhere in Europe and Japan.

The output will set out:

- Which healthcare professionals should be in which hospital, and when/if they should be moved between them.
- What central assets should be deployed to which hospitals, and when.
- How many patients of different categories should be sent to which hospital?
- How many respirators should be deployed where, and when should existing stocks be moved between hospitals if necessary.
- Optimiser will also account for the availability of consumables as a constraint – i.e. activities will only be planned if necessary equipment is available (for example oxygen, anaesthetic, PPE, consumables etc.)

The major constraints within this decision space are the limitations of a hospital and its capacity, as well as the medical logistic chain supporting it. There will be different hospital types across any given region, with varying capacities, capabilities and specialisms. Hospitals have a finite number of staff, resources and consumables. If it is impossible to find a solution with the resources available at the given hospital, then equipment and/or staff may need to move between hospitals - ideally these movement should be minimised. Alternatively Optimiser will, if necessary, direct patients to other hospitals - however this increases risk of infection.

Optimiser can be configured to focus on achieving one, or a weighted selection of, desired outcome(s). For the operational stage of crisis management, our understanding is that it would follow the rule of doing the most for the most, with the best use of finite resources. This is likely to be a mixture of maximum patient capacity, minimal movement of people and staff (to reduce cross contamination for example), to maximise patient recovery. These can be considered across one individual setting or group of multiple facilities.

Other planning functions that could be considered as a use case for Optimiser could be that of the requirement for mortuary services or testing centres, as well as the decommissioning of the field hospital facility once the demand on the nation's clinical infrastructure has been eased and the Defence Medical Service has been stood down.

Summary

We believe that this solution would bring benefit to work that the NHS and healthcare planners are performing during the current crisis - that could be exploited in future crisis, where there is the need to maximise the benefit from many possible outcomes within a constrained decision space.

There are many modelling and simulation systems currently available for healthcare planning that can simulate patient pathways through care. The coronavirus outbreak has also seen randomised clinical trials (RCT) of some highly-effective software emerging, involving AI and machine learning previously trialled in other areas of medicine, to develop personalised treatment plans for each patient and influence policies and collaboration. However, these are predicting what resources a patient will require to improve their chances of survival rather than the kind of multi-criteria decision-making of planning the resources for a whole hospital. Modelling and simulation predict the “outcome” from a set of inputs, and allow expert users to carry out what-if analysis - they do not tell a user how to best utilise what is available. They can tell you what will happen, not what you should do. Optimiser on the other hand uses modelling and simulation to explore multiple (many thousands of) options to find the “best” outcome (for a defined meaning of best). Optimiser tells you what to do to ensure the best outcome.

For these reasons we believe our approach is superior to, but complimented by, existing tools and approaches used in healthcare planning currently, and we are looking to identify potential partners in the NHS for future exploitation of the system.