Proceedings of the 2023 Winter Simulation Conference
C. G. Corlu, S. R. Hunter, H. Lam, B. S. Onggo, J. Shortle, and B. Biller, eds.

A GENERALIZED SYMBIOTIC SIMULATION MODEL OF AN EMERGENCY DEPARTMENT FOR REAL-TIME OPERATIONAL DECISION-MAKING

Alexander R. Heib
Christine S.M. Currie
Bhakti Stephan Onggo
Honora K. Smith

CORMSIS
University of Southampton
University Road
Southampton, SO17 1BJ, UK

James Kerr

Hampshire Hospitals NHS Foundation Trust
Aldermaston Road
Basingstoke, RG24 9NA, UK

ABSTRACT

We describe the design of a generalizable simulation model of an emergency department (ED) that forms part of a symbiotic simulation tool designed to improve short-term decision-making. While the paper will give an overview of the planned symbiotic simulation tool, our focus here is on the generalizability of the simulation model. The model is coded such that the routing logic of patient pathways are not explicitly defined but are instead included as an input parameter. By structuring the model this way, the pathways can instead be discovered through process mining methods on standard healthcare transactions data. This enables the simulation model to be applied to other EDs without redesigning all of the logical flows within the model. As symbiotic simulation tools are designed for ongoing use within the system they model, utilizing process mining also allows for automating recalibration of the patient pathways if changes occur in the physical system.

1 INTRODUCTION

The strain on hospital emergency departments (EDs) in the United Kingdom and Europe is well documented (Knight and Lasserson 2022). In late 2022, this became particularly pronounced when 1 in 7 patients in England waiting to be admitted into hospital from the ED had to wait over 12 hours (McCubbin and Triggle 2023) and fewer than 70% of patients were seen within 4 hours of arrival, the lowest rate since records began in 2004 (O'Dowd 2022). Alongside the overcrowding issues and limited resources, demand and service times in EDs typically have a high variance, which makes effective management difficult (Cildoz et al. 2018). The primary focus of much of the academic simulation research has been on designing so-called static models that are fitted to a particular ED and used to determine optimal operating strategies for any time of day. We have found little research that aims to optimize real-time (or near real-time) operational decisions based on the current state of the ED.

The aim of our research is the design of a generalizable simulation model of an ED that forms part of a symbiotic simulation tool designed to be installed within the department in oerder to support management in making operational decisions that need to be made daily. An example of such a decision is the allocation of staff to areas within the department.

A specific benefit of generalizable simulation models to real-time decision support systems is the ability to automate the updating of models based on new data. Updating the parameters used in arrival and service time distributions involves some statistical work, but no changes to the simulation logic. Conversely, with a typical discrete event simulation (DES) architecture, it is harder to change the pathways in the model without making fundamental changes to the logic. The design we introduce here allows that to happen much more easily, as we describe in Section 4.3.

The research project is being run in collaboration with Hampshire Hospitals NHS Trust who are responsible for two EDs in Winchester and Basingstoke, UK. Based on discussions with the management team, the two key relevant operating procedures that they can alter are the allocation of staff and patients to different areas of the ED. Based on data from Basingstoke, on approximately 80% of days between 150 and 200 patients attend the ED, with 10% of days having over 200 patients visit. Arrival rates exhibit seasonality over the year, within the week and within each day. Delays in moving patients out of the ED to inpatient wards within the hospital also add to the variability and the complexity of the system.

We turn next to a literature review (Section 2) before providing a description of the ED in Section 3. The design of out tool is described in Section 4 in which we describe the overall architecture, the process mining, some initial work on modeling demand and the simulation model implementation. Finally we conclude and describe our ideas for future work.

2 LITERATURE REVIEW

Our work draws on two key methodological concepts: symbiotic simulation and generalizable models. We discuss literature in these two areas before providing a review of simulation modeling of EDs.

2.1 Symbiotic Simulation

Aydt et al. (2008) define symbiotic simulation as a close association between a simulation system and a physical system, in the sense that the physical system provides data for the simulation system, and the simulation system output may control or influence decisions in the physical system. Oakley et al. (2020) and Onggo et al. (2020) found proposals of symbiotic simulation were introduced in the 1990s (Rogers and Gordon 1993; Davis 1998), although applications of the method were rare before 2005 (Oakley et al. 2020). Oakley et al. (2020) also found numerous applications of symbiotic simulation in industries such as manufacturing and transportation, with only a small number occurring in healthcare. Additionally, Harper and Mustafee (2019) found that despite there being research on short-term forecasting of ED demand, there is little research that utilises these forecasts for planning. Terms such as digital twin, real-time simulation, and online simulation have all been used synonymously with symbiotic simulation, but these terms also have a broader meaning that goes beyond what we are describing here.

While symbiotic simulation is more commonly applied to manufacturing problems because there is more routine use of automated data collection, there have been some applications in the health care domain. Of particular note, Harper and Mustafee (2019), Hoot et al. (2008), Espinoza et al. (2014), Marmor et al. (2009) and Tan et al. (2013) all model EDs, while Bahrani et al. (2013) and Oakley et al. (2020) include the ED in a model that describes more than just a single department, focusing on improving cardiac care and managing bed occupancy respectively throughout the hospital.

The DES model developed in Hoot et al. (2008) is designed to be initialized with real-time data with the aim of forecasting the state of the ED several hours into the future. The patient pathways were kept relatively simple, with treatment modeled as one process. The impact of bed blocking, where a patient is unable to be admitted to an inpatient ward from the ED because there are no inpatient beds available, is included under the the assumption that this would have an effect on capacity in the near future. Harper and Mustafee (2019) work on a similar problem, developing a hybrid model with the aim of reducing the likelihood of ED overcrowding. The model includes a forecasting unit to predict the occurrence of overcrowding in the near future, and a real-time DES unit to evaluate the performance of different mitigation strategies. Although the strategies to improve patient flow are not discussed in much depth in the paper, it is implied that a selection of designed strategies are compared, rather than optimizing over a general decision space. While this reduces the search space and so is likely to speed up computation, it relies on expert knowledge to design strategies, which may mean that superior strategies are left undiscovered.

Of the papers reviewed to date, we have not found any significant discussions about how to address the issue of evolving systems, particularly where there are changes to the system that affect the routing logic of entities. As symbiotic simulation tools are designed for ongoing use within the system they model, it is important that we consider that the system may change in a way that the model does not account for. By making a flexible simulation model, we reduce the chance of this happening.

2.2 Generalizable Models in Health Care

We use the definition given in Günal and Pidd (2011) that a generalizable model is generic and data-driven, so that it can be parameterized by input data to enable it to fit to different hospitals. This is similar to the definition in Boyle et al. (2022) that generalizable models "can represent multiple units and, through customization using real data, be used to investigate site-specific problems." One of the key benefits of a generalizable model is that it promotes model reuse. Our work extends that of Boyle et al. (2022) by taking a data-driven approach to determining the processes within the ED, using techniques from process-mining. The architecture of our simulation model makes it possible to do this rather than needing the simulation modeling expert to set up the logic for each new ED. Similar ideas were pursued in Mes and Bruens (2012) who built a flexible, generalizable model of an emergency post in the Netherlands.

2.3 Models of Emergency Departments

We performed a systematic review of operational research methods used for ED management. The search resulted in 126 papers from the years 2000 to 2020, of which 103 used a simulation in some capacity. Out of all the accepted papers in the review that used simulation methods, only two focused on short-term decision making, both of which used agent based methods to order patient priority (Ajmi et al. 2019; Bruballa et al. 2020).

A large number of these papers used a case study as the basis of the simulation model. They follow the standard modelling process of developing conceptual models of the department and then implementing this in the simulation. Because of this, these models often have specific features, reducing their potential for generalizability. De Boeck et al. (2019) also found that the modelling of specific units led to a reduction in model reuse.

One major difficulty in developing an ED model that can be applied to other EDs is the heterogeneity between departments. One example of these differences are only some include a separate paediatrics unit (Bair et al. 2010). Another example is whether the department is split into areas for treating different severity levels (Ghanes et al. 2015; Centeno et al. 2013) or different treatment needs (Al-Refaie et al. 2014), while some only have one queue for all patients (Bedoya-Valencia and Kirac 2016).

While it is possible to create a generic model that can capture the majority of behaviour in an ED, this comes at a cost of specificity. This may be acceptable for determining broad policies, but for a symbiotic simulation, where the decisions that can be made are on a much smaller scale, the simulation needs to be able make a distinction between a set of options. Thus we aim to make a simulation model that contains the detail of bespoke models, but generates this detail from process mining rather than conceptual modelling.

3 SYSTEM DESCRIPTION

In this section, we will explain the structure of the ED and the data that will be used in our model.

3.1 Structure of the Department

As the research is in collaboration with Hampshire Hospitals, the following description is specific to their EDs; however, this will contain similar pathways to many hospitals within the country, and internationally. The ED at Basingstoke and North Hampshire Hospital sees approximately 55,000 patients each year, meaning the department will treat on average 150 patients per day.

Patients arrive in the ED either by walking in or in an ambulance, with walk-ins making up 80% of attendances. The areas within the ED are: triage, Rapid Assessment and Treatment (RAT), paediatrics emergency unit, minor injuries, major injuries, resuscitation (resus) and the Emergency Decision Unit (EDU). The RAT is just for ambulance arrivals and the EDU is a place for patients if there is uncertainty in them needing to be admitted. Minor patients will sit in the waiting room in between treatments and tests while other patients will occupy a bed or a chair within their respective area. In addition to treatment, patients may need tests, either sample testing (e.g., blood tests) where the patient does not need to be present, or image testing (e.g., x-rays) where the patient needs to visit a diagnostic unit.

Around 80% of patients are discharged following a visit to the ED with the remaining 20% being admitted as inpatients. There can be a wait for an inpatient bed to become available and the decision to admit a patient is made as early as possible in a patient's treatment.

3.2 Data

We have hospital data on patient attributes such as age, entry type (walk in or ambulance), and admission decision, cross-referenced with timestamp data for arrival into the ED, and other activities such as attendances by physicians and nurse practitioners, any tests that were performed, when drugs were administered, and bed requests. These data can be used to estimate distributions for inter-arrival times and process flows through the department. Determining service time distributions requires some additional thought as time stamp data exists only for the start of activities and not their completion.

The transaction data also includes the name of the staff member carrying out the activity. This, combined with data on the total number of staff working in the ED on each shift, can be used to estimate the number of staff working in each area. A final source of real-time data relevant to the management of the ED is the bed occupancy in the rest of the hospital. This indicates the likely difficulty in finding an inpatient bed for patients being admitted.

4 SIMULATION MODEL DESIGN AND IMPLEMENTATION

In this section, we describe the different components of the symbiotic simulation tool and the simulation model implementation using Python. Figure 1 gives an overview of the architecture of the tool being built. The tool communicates with an external hospital data warehouse system. The three relevant datasets are (1) historical patient data that are used to calibrate and develop the models in the tool, (2) patient data that become available after the models are calibrated and (3) real-time patient data which is the most recent data that may come from sensors.

Our tool comprises calibration and synchronization units, a simulation model and an optimization model. Both decisions on staff allocation and patient length of stay in the ED occur on the scale of hours. Thus, to help with short-term decisions, we expect the symbiotic simulation tool to be run multiple times per day. Each time the tool is run, the synchronization unit is executed to give the simulation model an initial state based on the current state of the ED. This will be achieved using information on the patients and staff currently in the department. Additionally, for each run, we want to estimate the number and type of patients arriving in the near-future. We can determine these parameters based on past arrival data.

There are also many other parameters needed for the simulation model to function, but these are not situation-specific, and so do not need to be determined every time the symbiotic simulation tool is run. The calibration unit is composed of automated methods to determine these parameters. The idea is that it is executed once when installing the tool, and then infrequently, to ensure parameters are still accurately representing the ED.

We will firstly discuss how a flexible simulation model has been designed to allow for different ED setups. After this, we describe two components in the calibration unit, the seasonality analysis module and process discovery module, and how they become inputs into the simulation model.

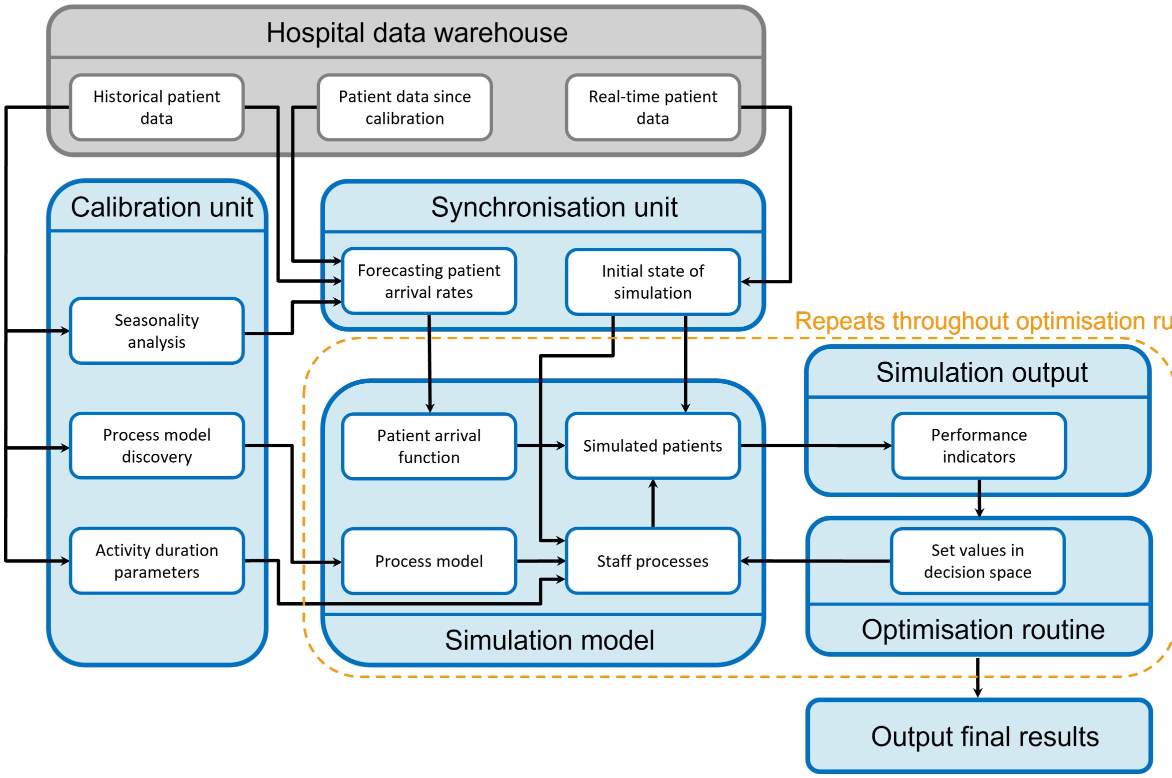

Figure 1: An overview of the planned project. The components in the calibration unit are intended to be run infrequently, at most every few months. The synchronization unit will run at the beginning of every simulation-optimization run.

4.1 Simulation Model Implementation

We developed the simulation model in Python using the SimPy package (https://simpy.readthedocs.io/en/latest/index.html). To make a generalizable ED model, instead of explicitly writing down the possible routes for patients to move through the ED, we defined a set of functions that affect a patient in different ways. These functions are equivalent to different tasks staff may perform for a patient and are the core of the simulation model. Examples of theses task functions include moving a patient from one area of the ED to the queue of another, attending to a patient, and admitting a patient to an area. The order and details of a patient's tasks are decided by the process model generated by the *process model discovery* module.

The modeling of the arrival process was also designed to be as generic as possible. Arrival occurs through a function instantiating patients in the simulation. One input of the function is a set of patient attributes, with no requirements on the number of attributes needed. This way, when initialising the model, the function can be called multiple times to initialize independent arrival processes, allowing patients with different attributes to have their own arrival processes. The purpose of this is explained in section 4.2.

The arrival process function also sets up the relevant objects needed for the simulation. Figure 2 shows the logic of the function. Similar to a typical simulation model, there is a timeout before the next patient arrival, at the end of which a patient is created and placed in a queue. In addition to this, a patient file keeping track of the patient information is created and placed in a store with all other patient files so any staff can access and update the file. Finally, the next task for the patient is determined through the use of the process model and placed in a store of all incomplete tasks.

Staff are modeled through a Python class. The class contains the task functions previously described, and an additional three functions to manage the performance of a task. The perform task function, visualised in Figure 3, retrieves the highest priority valid task for the staff member from the tasks store, calls the

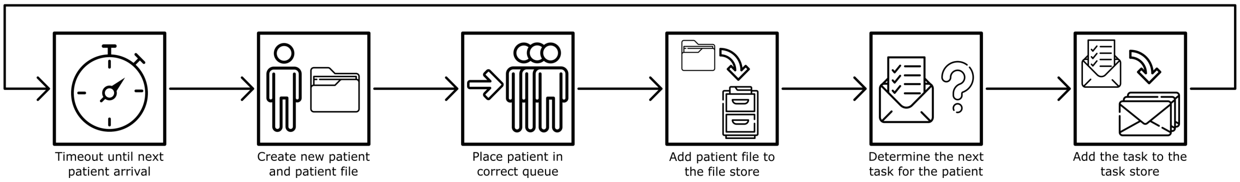

Figure 2: The arrival process of patients in the system.

appropriate task function, and assigns the next task associated with the patient, assuming the task was not a patient discharge. The next task is determined by finding the place of the patient in the process model and randomly selecting the next task if there is a choice. Additionally, for concurrent activities, multiple tasks can be placed in the task store for completion.

Sometimes the next task for the patient is not meant to be carried out immediately, for example a physician might need to wait a period of time after a patient is treated before being able to assess if the patient's health improves. Additionally, some tasks may not require a staff member, such as transferring an ambulatory patient from one area to another. To model these two features, a function handling automated tasks was defined. At the end of the perform task function, any tasks that do not need a staff member to complete are then passed immediately to the perform automated task. All other tasks are placed in the task store. The perform automated task function behaves in a similar way to the perform task function, with the exception that it is already assigned a task.

The final additional function in the staff class determines if a task in the task store that is available to the staff member is possible to complete. This function manages the admission of patients into different areas of the department. When a patient is placed in the queue of an area, an entry-type task is placed in the task store. However, there may not be a bed available to admit the patient. If the next task for a staff member to complete is an entry, but there is no space available, then they will continue to search for the next valid task. This function is then also able to handle more complex policies, such as allowing majors patients being admitted into a resus bed if there is capacity, and preventing staff from taking new patients if they need to be moved to a new area.

As the set of task functions have been written to be as generic as possible, so that they can apply to as many situations as possible, it does require matching the events found in the patient data with the correct function. This process can be simplified by creating a flow chart for classifying the behaviour of each activity. Additionally, any behaviour that cannot be replicated by the given functions can be included by writing a new function in the class.

For the simulation to run, it needs to be parameterised. One set of parameters is needed for the arrival process, and another is needed to determine the subsequent tasks for patients.

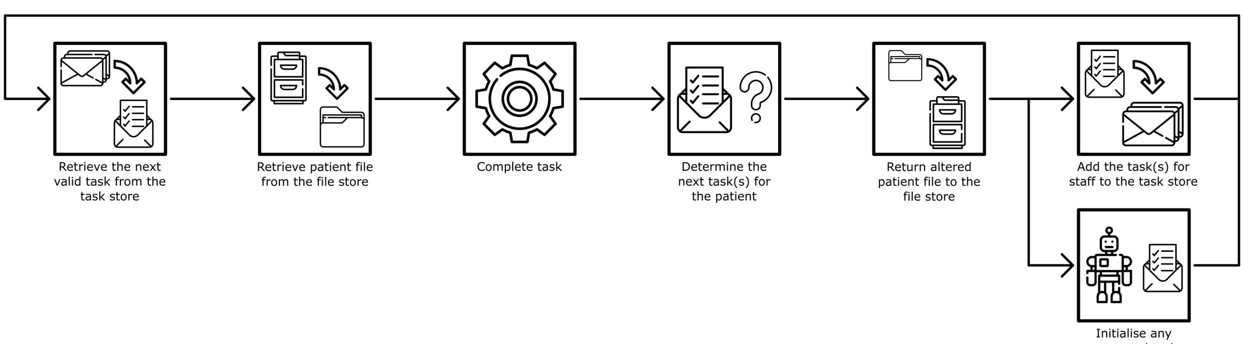

Figure 3: The process that staff in the system cycle through.

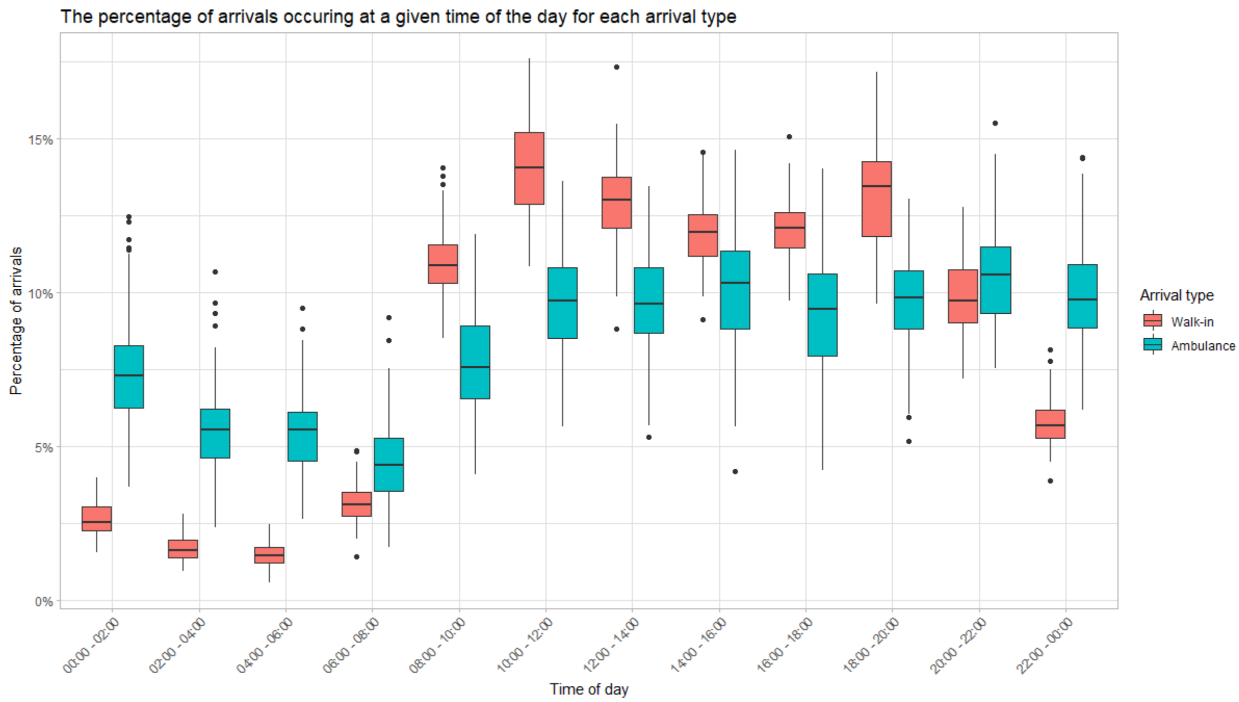

Figure 4: The percentage of arrivals occurring at each time of day for walk-in and ambulance patients.

4.2 Modeling Arrivals

The seasonality effects of patient arrivals are usually included in simulation models of EDs using non-homogeneous Poisson (NHPP) distributions. The modeling of seasonality often only affects the total number of patients arriving into the system, however there has been some research using more complex seasonal models. Borgman et al. (2015) modeled arrivals with a NHPP distribution, having a different distribution affecting the time of day for weekdays and weekends. Chen et al. (2020) and Cildoz et al. (2018) both modeled arrivals with a different NHPP distribution for each triage level. The papers mentioned only considered simple seasonality effects, though it is likely that there are far more complex interactions. From Figure 4, which uses data from Hampshire Hospitals ED, we can see that the arrival type has a seasonal pattern. Similar results can be found for other patient attributes, such as attendance type and age. To account for this, we have implemented an automated method to search for the significant seasonal effects in the patient arrival data. We also need to account for the relationship between patient attributes when determining seasonality, as the attributes are not independent; for example walk-in patients are more likely to attend minors than patients arriving by ambulance.

When using simulation for short-term decision-making about the best procedure to follow in the next few hours, it is important to capture all of the seasonality effects. This can be less important in static simulation modeling being used for more strategic, long-term decision-making.

We use a Poisson regression model, which can be written as a generalized linear model (GLM), to determine which interactions between seasonality variables and patient attributes need to be considered. We applied a stepwise regression search on multiple time periods to find interactions between variables. By partitioning the data, we were able to validate that similar sets of interactions were being identified. After this step, cross-validation was applied to the whole data set to eliminate any interactions that led to overfitting.

Preliminary results suggest that after taking into account interactions between patient attributes, there are a number of interaction terms between one seasonal variable and two patient attributes that need to be included. The time of day affected many patient attributes, with time of year having less influence. While the number of interactions can likely still be reduced for the simulation whilst retaining accuracy, there are still a number of significant seasonal effects. We anticipate the final set of results being used to set up routines for parameterization of the arrival processes for different patient types.

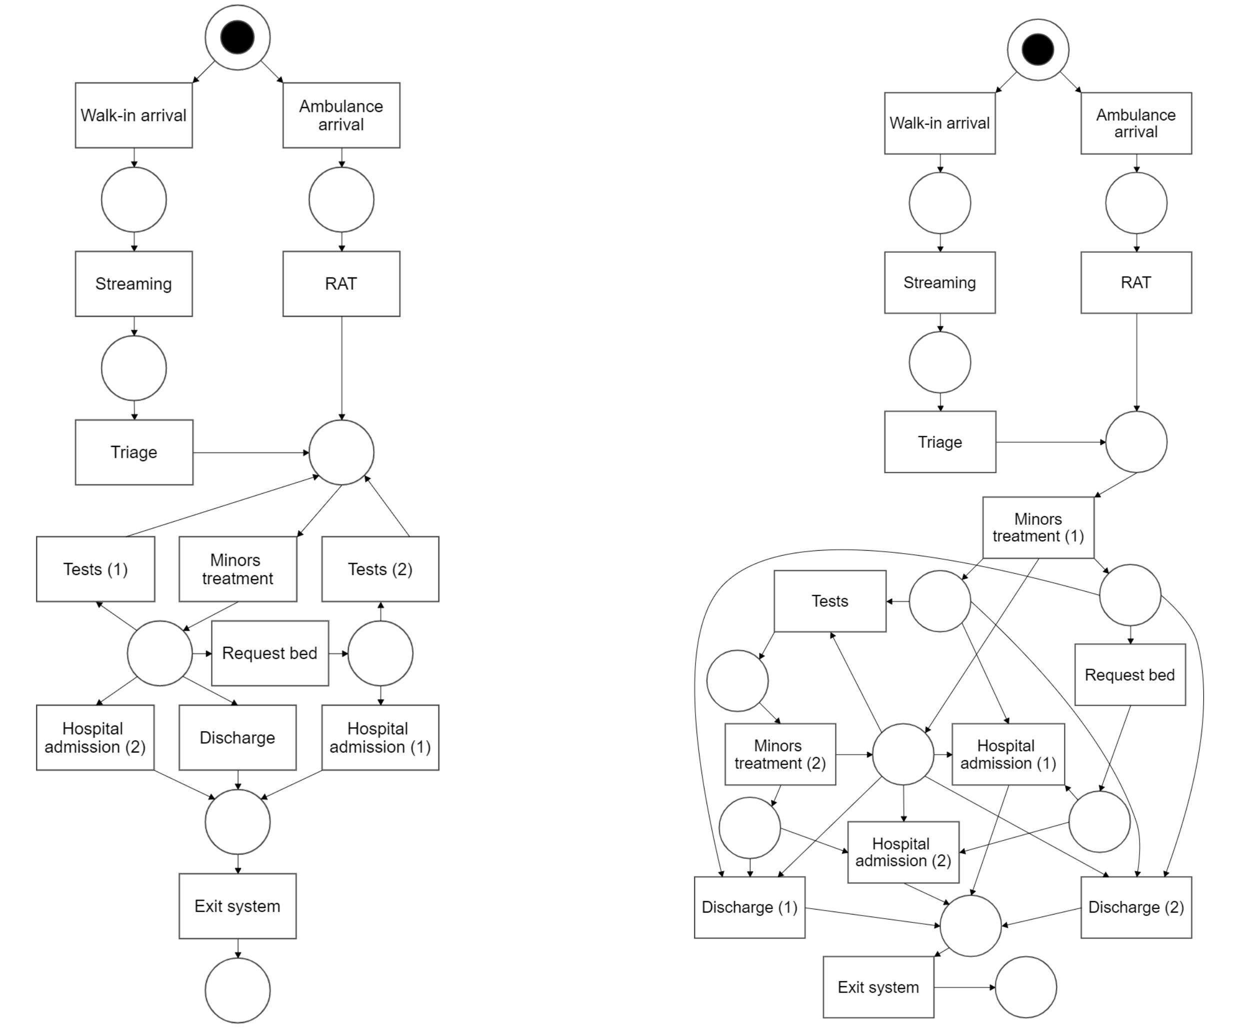

(a) A petri net derived from a transition system where states only contain the most recent activity.

(b) A petri net derived from a transition system where states contain all activities (unordered) in the partial trace.

Figure 5: Two petri net models derived from the same input data with different state functions applied.

4.3 Process Mining

We implement a process mining algorithm to find the care pathways that patients take through the ED. Then, instead of explicitly writing the discovered pathways in the simulation model, the output of the algorithm is fed directly into the model as a parameter. Each patient in the simulation will have an associated process model that will keep track of their progress, and also dictate the following activities.

Process mining involves the creation of process models through the application of algorithms on event logs (van der Aalst 2016). A requirement for applying process mining to a data set is for the event logs to contain unique identifiers for each case, i.e. a patient in the ED setting, an activity, and a timestamp.

These data are recorded in the hospital data warehouse, making process mining a good candidate for the modeling of patient pathways. Additionally, we sought a method such that relevant patient attributes, and potentially management choices can all be included in the process discovery.

There have been many process mining algorithms developed to discover process models from event logs, however many of these algorithms are rigid, in that there is little choice to influence the complexity of the resulting model (van der Aalst et al. 2010). This is a serious concern as it is possible to create either an overfitted model, which restricts behavior to pathways that have already been observed, or an underfitted model, which enables behavior that is distinctly different from what has been observed (Lugaresi and Matta 2021). A method that can result in different process models based on the level of abstraction can address this problem. This is relevant to the processes involved in an ED, where pathways may be heavily dependent on the history of the patient in the ED.

One method that can achieve this is a two step approach involving a state transition model and region-based mining (van der Aalst et al. 2010). The first step, making a state transition system, requires defining a state function. For each case, the state function is applied to each activity, and from this possible states and transitions can be discovered. The choice of state function is what affects the level of abstraction, for instance, a low abstraction state function may involve the exact order of previous activities that has occurred, whereas a high abstraction could be only the preceding activity. Hence, finding the level of abstraction needed to accurately model a system is a large component of this approach. The state transition system consists of states, which will be some representation of the process history for a case at some point in the case's journey through the process, and the transitions are activities. As an example, suppose the states are sets of the previous activities that have occurred. If a patient has been to triage and has been seen by a nurse practitioner in the minor injuries area then the state at this point is `[triage, minor injuries]`. If the next activity is having an x-ray, then the next state they will be in will be `[triage, minor injuries, x-ray]` and the transition between states is `x-ray`.

Once a state transition model is made, the second step is performed. This involves analyzing the state transition model to discover regions, which in turn relate to places in a petri-net process model (Cortadella et al. 1998). A region is defined as a subset of states where each activity adheres to exactly one behavior in respect to the region as listed (van der Aalst 2016):

1. All transitions for the activity direct a state outside of the region to a state inside the region.
2. All transitions for the activity direct a state in the region to a state outside of the region.
3. All transitions for the activity connect two states that are either both inside the region, or both outside the region.

Another common problem with process mining algorithms is an inability to discover either non-local dependencies, skippable activities, or both. A non-local dependency is where activity x will result in activity y being required, but not immediately following one another. This can certainly happen in an ED, for example a bed request will result in being admitted, but it is likely other activities will take place first. Skippable activities are also possible in an ED, for example what tests are performed will vary greatly for patients, it is not a case of having either an x-ray or a CT scan, some may have both, some may have neither. Both of these issues can be addressed with the two step method, further evidence that it is the best method for this project.

Figure 5 shows two petri net models that have been derived from the same data set but with two different state functions applied. Petri net models show when activities are allowed to happen. When all places (the circles) that directly precede an activity have a counter (the black circle), the activity is possible; the completion of an activity leads to all directly proceeding places to gain a counter.

Figure 5a is a model derived from a state function that generates states only from the most recent activity, that is, if a patient goes through activity a to activity b then there will be a state $\{a\}$ and state $\{b\}$ with a transition b going from the former state to the latter. Figure 5b is derived from a state function that generates states based on the complete unordered history of activities, that is, if a patient has gone through activities a twice, and b once, they will be in the state $\{a,a,b\}$, no matter what order they completed the activities.

Although Figure 5a is simpler, it allows for lots of incorrect behaviour, such as a patient being admitted to hospital without a bed request being made. This is because by only considering the previous activity, non-local dependencies are lost. Additionally, the model loses the potential for concurrency, despite that being possible with the tests and request bed activities. These problems are all fixed in Figure 5b, and give a much better representation of the process.

5 VERIFICATION AND VALIDATION

Both the simulation model and the process mining method have been verified to ensure they function as expected. For the simulation model, a simple process of patients arriving followed by a triage and then a treatment was set up to test the model functioned correctly. By examining the outputs of the model, we were able to verify visually that patients were being placed into the appropriate queues and exiting the system. As the simulation-optimization routine will need to have a small execution time, ideally under 20 minutes, another aspect of the simulation validation was ensuring that the model ran fast enough. One simulation run took on average 0.05 seconds to complete. The full model is expected to have approximately 5 times as many activities for each patient, and so an expected runtime of 0.5 - 1 second per simulation would lead to being able to execute at least 1000 simulations, even more with parallelization. For the process mining method, examples in the papers by Cortadella et al. (1998) and van der Aalst et al. (2010) were successfully replicated.

Validation methods will also be included in the final simulation tool, to ensure that the calibration unit is able to successfully model an ED. These methods will include statistical tests comparing simulated performance indicators to historical data, such as waiting times, queue lengths, and patient length of stay.

6 CONCLUSIONS AND FUTURE WORK

There is strong motivation to apply operational research methods to emergency departments, as they are complex systems that are often under pressure, with performance affecting the health and lives of patients. While there has been an abundance of research into improving emergency departments, this is often focused on supporting decision-making in the long-term (e.g. planning decisions, design decisions). Though this research is important, it does not always account for the large amount of variability that occurs in the short-term, especially important for real-time and near real-time decisions. This means that management are often left with little support to make confident decisions when dealing with the specific problems they encounter daily. Symbiotic simulation is a method that can support management making these decisions, as it is able to recreate the current state of the department, and evaluate a large number of potential options in the decision space.

It is common practice in the literature to build a new simulation model for each application, despite many departments being relatively similar. By utilising the flexibility of SimPy, and the method of automatic discovery of system structures with process mining, this research issue can be addressed. Once completed, this decision support tool should be applicable not only to the emergency departments at Hampshire Hospitals, but any emergency department in the UK with only minor modifications.

Another reason for constructing the model this way, is that it has a greater chance at remaining a reliable tool in the future. By allowing for the routing logic of the simulation to be recalibrated, the simulation will still be able to model the department after structural changes occur. It would also benefit if data collection improves, as if more events for a patient are recorded digitally, then access to this will lead to a more accurate process model.

The immediate next step on the project will be to fully parameterize the simulation model for our collaborating hospital and evaluate its performance by comparing with historical data. We also aim to improve the process mining algorithm to better handle systems with skippable activities, as the current method becomes increasingly slow for larger problems. Methodologically, the final jigsaw piece in the symbiotic simulation tool we are building is the optimization algorithm, which will enable support of real-time decision-making by the management team.

ACKNOWLEDGMENTS

Alex Heib was supported by EPSRC grant number EP/T517859/1. We are also grateful to the Hampshire Hospitals NHS Foundation Trust for providing data and expertise.

REFERENCES

Ajmi, F., H. Zgaya, S. Ben Othman, and S. Hammadi. 2019. "Agent-based Dynamic Optimization for Managing the Workflow of the Patient's Pathway". *Simulation Modelling Practice and Theory* 96.

Al-Refaie, A., R. H. Fouad, M.-H. Li, and M. Shurrab. 2014. "Applying Simulation and DEA to Improve Performance of Emergency Department in a Jordanian Hospital". *Simulation Modelling Practice and Theory* 41:59–72.

Aydt, H., S. J. Turner, W. Cai, and M. Y. H. Low. 2008. "Symbiotic Simulation Systems: An Extended Definition Motivated by Symbiosis in Biology". In *2008 22nd Workshop on Principles of Advanced and Distributed Simulation*, 109–116.

Bahrani, S., R. B. Tchemeube, A. Mouttham, and D. Amyot. 2013. "Real-Time Simulations to Support Operational Decision Making in Healthcare". In *Proceedings of the 2013 Summer Computer Simulation Conference*, SCSC '13. Vista, CA: Society for Modeling & Simulation International.

Bair, A. E., W. T. Song, Y.-C. Chen, and B. A. Morris. 2010. "The Impact of Inpatient Boarding on ED Efficiency: A Discrete-Event Simulation Study". *Journal of Medical Systems* 34(5):919–929.

Bedoya-Valencia, L., and E. Kirac. 2016. "Evaluating Alternative Resource Allocation in an Emergency Department Using Discrete Event Simulation". *Simulation-Transactions of the Society for Modeling and Simulation International* 92(12):1041–1051.

Borgman, N. J., M. R. K. Mes, I. M. H. Vliegen, and E. W. Hans. 2015. "Improving the Design and Operation of an Integrated Emergency Post via Simulation". *Journal of Simulation* 9(2):99–110.

Boyle, L. M., A. H. Marshall, and M. Mackay. 2022. "A Framework for Developing Generalisable Discrete Event Simulation Models of Hospital Emergency Departments". *European Journal of Operational Research* 302(1):337–347.

Bruballa, E., A. Wong, D. Rexachs, and E. Luque. 2020. "An Intelligent Scheduling of Non-Critical Patients Admission for Emergency Department". *Ieee Access* 8:9209–9220.

Centeno, A. P., R. Martin, and R. Sweeney. 2013. "REDSim: A Spatial Agent-Based Simulation for Studying Emergency Departments". In *Proceedings of the 2013 Winter Simulation Conference*, edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. R. Hill, and M. E. Kuhl, 1431–1442. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Chen, W., H. Guo, and K.-L. Tsui. 2020. "A New Medical Staff Allocation via Simulation Optimisation for an Emergency Department in Hong Kong". *International Journal of Production Research* 58(19):6004–6023.

Cildoz, M., F. Mallor, and A. Ibarra. 2018. "Analysing the ED Patient Flow Management Problem by using Accumulating Priority Queues and Simulation-Based Optimization". In *Proceedings of the 2018 Winter Simulation Conference*, edited by M. Rabe, A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, 2107–2118. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Cortadella, J., M. Kishinevsky, L. Lavagno, and A. Yakovlev. 1998. "Deriving Petri Nets from Finite Transition Systems". *IEEE Transactions on Computers* 47(8):859–882.

Davis, W. 1998. *Handbook of Simulation : Principles, Methodology, Advances, Applications, and Practice*. New York Norcross, Ga: Wiley Co-published by Engineering & Management Press.

De Boeck, K., R. Carmen, and N. Vandaele. 2019. "Needy Boarding Patients in Emergency Departments: An Exploratory Case Study Using Discrete-Event Simulation". *Operations Research for Health Care* 21:19–31.

Espinoza, C., J. Pascual, F. Ramis, D. Bórquez, and J. A. Sepúlveda. 2014. "Real-Time Simulation as a way to Improve Daily Operations in an Emergency Room". In *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 1445–1456. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Ghanes, K., M. Wargon, O. Jouini, Z. Jemai, A. Diakogiannis, R. Hellmann, V. Thomas, and G. Koole. 2015. "Simulation-Based Optimization of Staffing Levels in an Emergency Department". *Simulation-Transactions of the Society for Modeling and Simulation International* 91(10):942–953.

Günal, M. M., and M. Pidd. 2011. "DGHPSIM: Generic Simulation of Hospital Performance". *ACM Transactions on Modeling and Computer Simulation* 21(4).

Harper, A., and N. Mustafee. 2019. "A Hybrid Modelling Approach Using Forecasting and Real-Time Simulation to Prevent Emergency Department Overcrowding". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, K.-H. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, 1208–1219. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Hoot, N. R., L. J. LeBlanc, I. Jones, S. R. Levin, C. Zhou, C. S. Gadd, and D. Aronsky. 2008. "Forecasting Emergency Department Crowding: A Discrete Event Simulation". *Annals of Emergency Medicine* 52(2):116 – 125.

Knight, T., and D. Lasserson. 2022. "Hospital at Home for Acute Medical Illness: The 21st Century Acute Medical Unit for a Changing Population". *Journal of Internal Medicine* 291(4):438–457.

Lugaresi, G., and A. Matta. 2021. "Automated Manufacturing System Discovery and Digital Twin Generation". *Journal of Manufacturing Systems* 59:51–66.

Marmor, Y. N., S. Wasserkrug, S. Zeltyn, Y. Mesika, O. Greenshpan, B. Carmeli, A. Shtub, and A. Mandelbaum. 2009. "Toward Simulation-Based Real-Time Decision-Support Systems for Emergency Departments". In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, 2042–2053. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

McCubbin, J., and N. Triggle. 2023, Jan. "Ambulances Taking 90 Minutes to get to 999 Calls". *BBC News*. https://www.bbc.co.uk/news/health-64254249, accessed 12th October 2023.

Mes, M., and M. Bruens. 2012. "A Generalized Simulation Model of an Integrated Emergency Post". In *Proceedings of the 2012 Winter Simulation Conference*, edited by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. M. Uhrmacher, 872–882. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Oakley, D., B. Onggo, and D. Worthington. 2020. "Symbiotic Simulation for the Operational Management of Inpatient Beds: Model Development and Validation using Delta-Method". *Health Care Management Science* 23:153–169.

Onggo, B. S., C. G. Corlu, A. A. Juan, T. Monks, and R. de la Torre. 2020. "Combining Symbiotic Simulation Systems with Enterprise Data Storage Systems for Real-Time Decision-Making". *Enterprise Information Systems*.

O'Dowd, A. 2022. "NHS in England: Four Hour Emergency Department Target Hits Record Low". *The BMJ* 379(o2719).

Rogers, P., and R. Gordon. 1993. "Simulation for Real-Time Decision Making in Manufacturing Systems". In *Proceedings of 1993 Winter Simulation Conference*, edited by G. Evans, M. Mollaghasemi, E. Russell, and W. Biles, 866–874. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Tan, K. W., W. H. Tan, and H. C. Lau. 2013. "Improving Patient Length-of-Stay in Emergency Department Through Dynamic Resource Allocation Policies". In *2013 IEEE International Conference on Automation Science and Engineering (CASE)*, 984–989.

van der Aalst, W. 2016. *Process mining: Data Science in Action*. 2nd ed. Berlin: Springer.

van der Aalst, W. M., V. Rubin, H. M. Verbeek, B. F. van Dongen, E. Kindler, and C. W. Günther. 2010. "Process mining: A Two-Step Approach to Balance Between Underfitting and Overfitting". *Software & Systems Modeling* 9(1):87–111.

AUTHOR BIOGRAPHIES

ALEXANDER R. HEIB is a PhD Student at the University of Southampton within the Operational Research group in Mathematical Sciences. His research focus is simulation modeling in healthcare settings. His email address is arh1n19@soton.ac.uk.

CHRISTINE S.M. CURRIE is a Professor of Operational Research in Mathematical Sciences at the University of Southampton and a member of the Centre for Operational Research, Management Sciences and Information Systems (CORMSIS). She is Editor-in-Chief for the Journal of Simulation. Her research interests include simulation optimization, applications of simulation in health care, optimal pricing and disaster management. Her email address is christine.currie@soton.ac.uk and her homepage is https://www.southampton.ac.uk/people/5wzzxf/professor-christine-currie.

BHAKTI STEPHAN ONGGO is a Professor of Business Analytics at the University of Southampton. He is a member of the Centre for Operational Research Management Sciences and Information Systems (CORMSIS) and Centre for Healthcare Analytics. His research interests include simulation modeling methodology and its applications in health care, disaster management and supply chain. His e-mail address is b.s.s.onggo@soton.ac.uk. His website is https://bsonggo.wordpress.com/.

HONORA K. SMITH is a Visiting Fellow in Mathematical Sciences at the University of Southampton and member of the Centre for Operational Research, Management Sciences and Information Systems (CORMSIS). Now retired as a Lecturer in Operational Research (OR), she continues with her research interest in the applications of OR modeling in healthcare. Her e-mail address is Honora.Smith@soton.ac.uk. Her website is https://www.southampton.ac.uk/maths/about/staff/hks1u06.page.

JAMES KERR is a Consultant in Emergency Medicine at Hampshire Hospitals NHS Foundation Trust and the Associate Medical Director for Clinical Strategy. He also works for NHS England as a Regional Clinical Advisor for Urgent and Emergency Care. He received his MBA with distinction from Durham Business School and is interested in education and quality improvement. His e-mail address is james.kerr@hhft.nhs.uk and his website is https://about.me/jameskerr.