

HYBRID SIMULATION FOR HEALTH AND SOCIAL CARE: THE WAY FORWARD, OR MORE TROUBLE THAN IT'S WORTH?

Sally C. Brailsford
Joe Viana
Stuart Rossiter
Andrew (Amos) Channon
Andrew J. Lotery

University of Southampton
Southampton SO17 1BJ, UK

ABSTRACT

This paper describes the process of developing a hybrid simulation model for a disease called age-related macular degeneration (AMD), a common cause of sight loss in people aged over 65. The model is implemented in the software AnyLogic, and combines discrete-event and agent-based simulation. Embedded in each agent there is also an individual compartmental model for disease progression. The overall aim of the hybrid model was to use the specific example of AMD to explore the wider links between the health and social care systems in the UK. We discuss the challenges of model development and the rationale for our modelling decisions, and reflect upon the advantages and disadvantages of using a hybrid model in this case.

1 INTRODUCTION

The Care Life Cycle project (CLC 2013; Brailsford et al. 2011) is a five year multidisciplinary research program at the University of Southampton, funded by the UK Engineering and Physical Sciences Research Council (EPSRC). The project, which began in 2010, is one of four major EPSRC-funded projects whose overarching aim is to apply “complexity science” methods to real-world problems: caring for an ageing population, the resilience of industrial ecosystems, food supply chain safety and immigration and diversity. The Care Life Cycle (CLC) focuses on issues of supply and demand for health and social care in a changing and ageing society. It is led by a demographer, Jane Falkingham, and involves computer scientists, social statisticians, gerontologists and operations researchers. The CLC uses both complexity science and traditional simulation modeling methods, and its aim is not to build a massive mega-model of the whole UK health and social care system, but rather to develop a suite of linked models which each tackle different aspects of this enormously complex system. These models operate at different levels of resolution, some individual-based and some at population level.

At the 2012 Winter Simulation Conference, three of these models were presented in an invited session (Noble et al. 2012; Brailsford et al. 2012; Viana et al. 2012). The third of these was a hybrid simulation model, developed in the software AnyLogic, for the eye disease age-related macular degeneration (AMD). The model combines discrete-event and agent-based simulation, and also, embedded in each agent, a pair of individual compartmental models (one for each eye) for disease progression. The overall aim of the hybrid model was to use the specific example of AMD to explore the wider links between the health and social care systems in the UK. In this paper we discuss the rationale for our modelling decisions, the challenges of model development, and reflect upon the advantages and disadvantages of using a hybrid model in this case.

2 BACKGROUND

In England and Wales, almost all health care is funded by taxation and provided free at the point of delivery by the National Health Service (NHS). At a local level, health care services – surgical operations, medicines, treatments etc - are purchased from NHS hospitals and primary care providers on behalf of their local population by a group of clinicians (a Clinical Commissioning Group, CCG). Each CCG has the challenging task of deciding how best to use their allocated budget. In contrast, social care (ranging from help with “activities of daily living” such as shopping, cooking, using the bathroom and so on, through to residential care) is mostly paid for by the individual themselves. Only the very poorest people with critical social care needs receive financial support from the state, and this is means-tested. The majority of people either purchase care directly from private providers, or receive “informal” care, i.e. from family, neighbors, or the charitable sector. State-funded care is not provided nationally but at local level by the Local Authority (county council) where the person resides. Health and social care are therefore paid for by different organizations from different budgets, and this can lead to complex relationships between the health and social care systems, especially in the elderly who are the greatest consumers of both.

Age-related macular degeneration (AMD) is a condition which exemplifies some of the difficulties this can lead to. AMD has two forms, wet and dry, and is the leading cause of sight loss in people aged over 65. It causes loss of central vision, meaning that people can no longer drive, read, watch TV, recognize faces, or even move safely around their homes. AMD therefore severely impacts on the quality of life of elderly people and leads to loss of independence. In the past there was no treatment for either form of AMD and people with the condition would gradually go blind, leading to increased need for social care and an increased burden on Local Authority budgets. However, in recent years new treatment has become available for wet AMD, delaying or even preventing deterioration in sight. This treatment requires regular monthly injections into the eyes, which have to be administered by a specialist clinician in a hospital outpatient setting. This, in combination with the higher volume of AMD patients due simply to the ageing population, has led to greatly increased demand for clinic appointments. Wait times in the clinic were rising, and some patients were even leaving without being treated because their transport had arrived to take them home again. Many older, less mobile patients are reliant on hospital transport, as they are unable to use public transport and have no local family support.

3 CHRONOLOGY: HOW THE MODEL CAME ABOUT

Initially, the Centre for Operational Research, Management Science and Information Systems (CORMSIS) at the University of Southampton was contacted by Dr Steve Lash, a consultant ophthalmologist in the Eye Unit at University Hospital Southampton. Dr Lash wondered whether simulation could be used to improve patient flows in the Eye Unit, relating in particular to scheduling surgeries for a condition called vitreoretinal detachment. This led to some joint work between CORMSIS and the Eye Unit, during the course of which we discussed the Care Life Cycle project and the possibility of further collaboration on eye conditions affecting older people, where there might be interactions between health and social care. Dr Lash suggested that we spoke to his colleague Professor Andrew Lotery, who specializes in age-related conditions such as AMD, glaucoma, cataract and diabetic retinopathy. After initial discussions we decided to focus on AMD, since it was of particular interest to Professor Lotery as he had been involved in clinical trials of the new treatment. Moreover, we all felt that modeling multiple conditions would have been a step too far in terms of complexity.

Professor Lotery described the problems in the outpatient clinic: long waits and untreated patients. At face value, these problems appeared to call for a traditional discrete-event simulation (DES) approach. The ophthalmology outpatient clinic is a complex queuing network where patients pass through a series of resource-constrained processes, such as booking in, eye examinations, complex interventions and specific treatments and so on, each requiring different rooms, types of equipment and/or staff. The system is complex and stochastic, since patient arrivals and activity durations are all variable and the clinic does not

only deal with AMD but many other conditions. Could these processes be redesigned to make better and more efficient use of the available resources, improve patient outcomes and eliminate the problems of long waits and untreated patients, or was it actually impossible to deal with the demand without additional resources? If the latter turned out to be the case, then the model would provide strong evidence to back up a request to the hospital management for additional resources.

There are countless examples of DES models of hospital clinics in the literature (Jun et al. 1999; Brailsford et al. 2009) and such a model, although it would be challenging and very time-consuming to develop in terms of data collection, would be a standard application of DES. As such it did not present a great deal of research interest, although it would have been of direct benefit to patients and staff in the Eye Unit. Although one aim of the CLC project is to work closely with local and national stakeholders, policy-makers and planners in order to make a real-world impact and improve the delivery of care, essentially the CLC is funded by the EPSRC to undertake innovative research in modeling methodology. Therefore its principal aim is to be a breeding ground for new modeling approaches. Therefore, we felt that developing yet another conventional DES clinic model was not something which our funders would regard as desirable.

Moreover, of course in reality the ophthalmology clinic does not operate in isolation. One of the great challenges of modeling hospitals, indeed any healthcare system, is the difficulty of drawing model boundaries, because in healthcare, “everything affects everything else”. Both health and social care interventions could affect the rate of sight loss in AMD patients. An example of a social care intervention might be to provide personal taxi transport for all AMD patients, which would reduce the number of no-shows and prevent patients who did attend from having to leave without being treated. We therefore decided to use AMD as an exemplar of the connections between the health and social care systems. By extending the clinic DES model to include some of these connections, we could provide an excellent illustration of the role of modeling to understand some of these wider system effects. We therefore widened the scope of the model to include social care within the community as well as medical treatment in the hospital clinic. From a technical perspective, this would also enable us to explore hybrid modeling approaches, which are discussed in section 5.

Our model has numerous stakeholders and serves several purposes. For the Eye Unit, the model serves the traditional “what-if” purpose of evaluating the effects of different options for providing additional resources, extending clinic opening hours, etc. However for the wider NHS, and for the local authority, the model provides a risk-free (and politics-free) environment for learning about the system, exploring a wide range of potential interventions and understanding the connections between them. By allowing the NHS and local authorities to see the whole system, and the long-term effects in one sector of short-term interventions in the other sector, the model can help break down organizational barriers. Of course, the same principles could be applied to many other long term conditions.

The wider system model also served a purpose for the CLC team: it provided an opportunity to explore different combinations of simulation approaches. Arguably, the choice of a hybrid model, as compared with other “whole-system” modeling approaches, was partly driven by academic curiosity. However, a hybrid approach enabled us to capture different aspects of the real-world problem that could not have easily been captured by any single modeling paradigm. Being part of a funded research project gave us the opportunity, the time and the space to develop a more innovative model which would be of greater research interest to the other CLC academic investigators, the funder EPSRC, and the wider academic community.

4 CONTEXT: LITERATURE AND SOFTWARE

Brailsford, Desai and Viana (2010) summarize the literature on the use of hybrid simulation models in health and social care, focusing on DES and system dynamics (SD), and present two examples of hybrid models. However, they argue that the “holy grail” of genuinely combining the philosophy of both approaches in the field of Operations Research has not yet been attained. Despite a long history of model-

ing tools in other disciplines such as engineering or computer science, which combine continuous and discrete parameters, there is as yet no genuinely integrated simulation methodology which reflects the differing world-views of DES and SD modelers. Morecroft and Robinson (2006) contrast these world-views, and illustrate the use of each approach for modeling the same system (fishery stocks). Tako and Robinson (2009) compare the differences in model-building approaches by users of DES and SD. Chahal and Eldabi (2008) identify three modes in which DES and SD can be combined in practice, which they term *hierarchical*, *process environment*, and *integrated*.

Most modern DES software tools contain some facility to model continuous, as well as discrete event phenomena, and can therefore be adapted to depict the underlying structures of SD models. Similarly, most SD software tools contain the facility to model queuing processes or stochastic variables. Nevertheless, these packages are essentially either a DES environment with some continuous features, or an SD environment with some discrete or stochastic features. The software AnyLogic (www.xjtek.com/AnyLogic) still remains the only package which can genuinely represent both DES and SD models (and also agent-based models). It offers nearly all the main features available in paradigm-specific software tools such as Vensim, Powersim and iThink (SD tools) and Arena, Witness, Simul8 and ProModel (DES tools). However, this flexibility comes at a price. The user needs some familiarity with Java code, and there is a “jack-of-all-trades” effect: an expert user of any of the above paradigm-specific tools may at times yearn to use their preferred software.

Viana (2011) describes a hybrid DES-SD model for the sexually transmitted infection (STI) chlamydia. This has several aspects in common with the AMD model, although it does not reflect social care. A DES model of the hospital STI clinic is embedded within an SD model of the infection process within the wider community. People do not book appointments for the STI clinic, but just show up, are tested, and if necessary treated with a simple course of antibiotics. If the clinic is crowded, the queues get long and people are unwilling to wait. Therefore they do not get tested (or treated). This then increases the proportion of infected people in the community, leading to an increase in new infections and further increases the demand for the clinic: the SD part of the model captures this feedback effect. However, this model uses two separate specialist packages, Simul8 and Vensim, which communicate via an Excel interface.

5 MODEL DESCRIPTION

The model structure is described in Viana et al. (2012). Essentially, it combines an agent-based model representing individuals with AMD with a DES model of the outpatient clinic. The patient agents contain two simple embedded state transition models which represents the progression of AMD in each eye. This is affected by treatment, which slows down the disease process. Each patient has a social care need status, which corresponds to the categories of need used by UK Local Authorities: none, low, moderate, substantial and critical (Department of Health, 2003) and this, in conjunction with the level of social care provision, affects their probability of clinic attendance. Social care provision is represented by a state chart which consists of three states: not required, partly met and fully met. The agents interact with the DES model when the scheduled time of their clinic appointment arrives. The patient may (or may not) attend the clinic, and may (or may not) receive treatment, depending on the congestion in the clinic and the overall performance of the clinic.

The model has been parameterized with data from a combination of different sources, including the English Longitudinal Study of Ageing (ELSA 2012), the Age-Related Eye Disease Study Group (Age-Related Eye Disease Study Research 2000; Age-Related Eye Disease Study Research 2005), and a study commissioned by the Eye Unit (Access Economics PTY Limited 2009). Other information, relating to the operation and performance of the Eye Unit, was obtained directly from the Unit staff. Further details can be found on the CLC website (CLC 2013) at <http://www.soton.ac.uk/clc/publications/supplementary>, together with the parameters used from these data sources in different parts of the model. Figure 1 depicts the clinic layout and illustrates the AnyLogic DES environment.

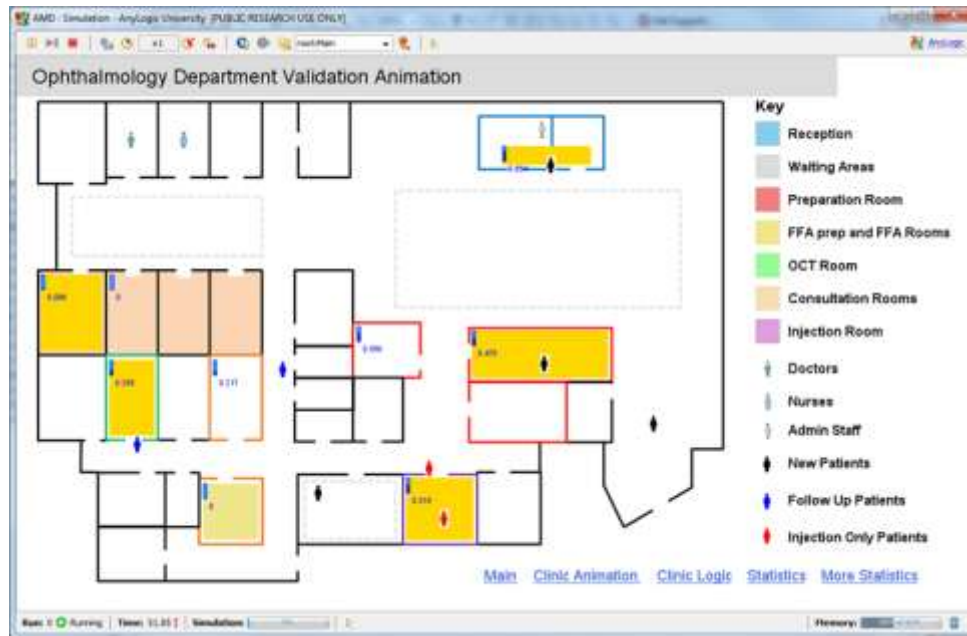


Figure 1. The Eye Unit outpatient clinic: DES model layout

6 MODEL DEVELOPMENT PROCESS

The main aim of this paper is to describe the hybrid modelling process, from conceptualization to coding, examining the key modelling decision points and discussing the rationale for each choice. We shall not focus on the detailed technicalities of AnyLogic, as this is not intended to be a paper about model coding, other than to discuss the points at which decisions or compromises had to be made. Furthermore, the paper is not intended to be a systematic critique of AnyLogic but just reflects our individual experience. The two people who coded the model have different disciplinary backgrounds. Joe Viana is an Operations Researcher: he is an experienced and expert user of Simul8 and Vensim, but he had no previous experience in using AnyLogic or Java. Stuart Rossiter is a computer scientist and an expert programmer, with experience of several agent-based simulation frameworks (MASON and Repast Symphony).

As discussed in Section 3, the model served a number of different purposes which drove the decision to use a hybrid approach. Had we been undertaking the modeling on behalf of a single stakeholder, the Eye Unit (for example, if we were a business consultancy and the Eye Unit were our client), we would almost certainly have developed a standalone DES model, and we would almost certainly not have used AnyLogic but would have chosen a DES tool with which we were familiar. In fact, AnyLogic allows the user to turn off the different components of a model, and so it is possible to use our model as a straightforward DES of the clinic if required.

One possibility would have been to use the same approach as Viana's chlamydia model (2011), namely to combine the DES clinic model (coded in Simul8) with a system dynamics model (coded in Vensim) of the social care process. However, there were a number of arguments against this. Firstly, it was by no means obvious that SD was the ideal choice for the wider population model. Chlamydia is an infectious disease and the feedback effects between infected and susceptible populations can be captured very naturally and easily in SD. However, AMD is not infectious: its progression occurs inside each individual's eyes and is affected by that person's treatment. Moreover the aspects of social care need that we wished to model also varied across individuals. Therefore, we needed a modeling approach which could capture individual behavior as well as the clinical progression of disease inside that individual. Sec-

only, from a theoretical perspective the conceptual approach would not be all that original as it would be replicating, albeit for a totally different disease, the approach used in the chlamydia model.

We could possibly have used DES for the whole model and just extended the scope of the model, but a hybrid simulation was a more natural fit. The social care side is fundamentally concerned with individual interactions (suggesting a need for agent-based modeling), whereas the clinic is fundamentally about workflows (hence DES). Sight loss is a continuous process and this naturally suits a differential equation or SD approach. However, individual decisions, especially of clinicians, could significantly affect outcomes, and hence we were interested in the possibility of including clinic staff as decision-making agents. An agent-based approach would also allow us to capture interactions between older people and their carers. Moreover, another CLC project was under way, using agent-based modeling to represent demographic change and family formation, with explicit interactions between care-giving and care-receiving agents (Noble et al. 2012). This population dynamics model was being coded from scratch in the programming language Python. Whilst we also believe that the interaction between individuals is the fundamental mechanism for modeling social care, we also wanted to compare an individual-based approach where individuals operate as largely independent “statistical individuals” (Courgeau, 2007), transitioning between care states and AMD stages according to population-level statistical models. This is the essence of the microsimulation/microanalytical simulation approach (see §4 of Gilbert & Troitzsch (2005)), and is effectively an agent-based model without the AI-oriented aspects of agency discussed by researchers such as Wooldridge (2002). Currently, the only agents in the AMD model are individuals with AMD, and we do not explicitly model their families or other support mechanisms.

The CLC had recently acquired AnyLogic and so this seemed the ideal opportunity to test it out on a real project. Therefore the decision was taken to use AnyLogic to build the model and to use an agent-based approach for the social care aspects. This also gave us the potential to include geographical location as one of the agents’ characteristics, although it is not currently included in the model. As it stands, the agents are distributed at random on a rectangular space, and distance from the hospital is only modeled implicitly as part of the probability of attendance. We plan to include this feature in due course, as it will enable us to represent where patients live in relation to the hospital, and investigate scenarios such as the use of a mobile clinic which could visit community centers and residential care homes. Because of our desire to compare this model with other styles of model, we also wanted to construct the model in as modular and loosely-coupled manner as possible, so that we could potentially ‘swap-in’ and ‘swap-out’ different sub-components, rather than focus solely on parametric variation of a single model. For the agent-based aspects, this also gave us the chance to compare the model development process in a commercial tool (AnyLogic) with development from scratch in the Python-based population dynamics model. Perhaps more interestingly from the hybridization perspective, we also envisage that we could couple models of differing paradigms together at a more granular level: for example, clinicians could themselves be agents with differing characteristics (skill level, tiredness, etc.) who make independent decisions and interact with the rest of the clinic (e.g., in signing off junior clinician diagnoses, or in adjusting their appointment timetables due to increasing demand or problems fitting patients into the appointment slots used previously).

As part of this natural experiment in model-building, we wanted to use as much of the default AnyLogic functionality as possible. Joe began by coding the DES clinic model. This had two purposes: firstly, he was far more familiar with the DES paradigm and so this was a good way to learn a new package and acquire the necessary knowledge of Java. Secondly, Professor Lotery and the Eye Unit were naturally more engaged with the operations of the AMD clinic. Joe adopted a standard approach for building a DES model: talk to the stakeholders to understand the system, visit the clinic to see for himself, jointly decide on the model scope, jointly map out the processes, understand the resource requirements, and identify and collect data to populate the model from the clinic itself where possible. The visual version (Figure 1) was developed for validation purposes. For example, the model animation showed queues of patients building up under the same circumstances that they do in real life: this face validity, combined with

the realistic layout of the clinic on the screen, was very helpful in gaining acceptance from the clinic staff. Unsurprisingly, the animation was a powerful aid to communication and “stakeholder buy-in”: it greatly impressed staff in the Eye Unit.

At the same time, Joe began to develop an initial version of the agent-based section of the model and explore options for linking this to the clinic DES. System dynamics models were embedded in each agent to represent the person’s sight level. Sight is modeled as a percentage of perfect (20-20) vision which deteriorates over time if the AMD is untreated: it also naturally deteriorates with age. The parameters are derived from the clinical literature as described in Section 5. Each agent also contained a very simple DES model which scheduled the time of the agent’s next clinic appointment and then transferred the agent into an entity arriving at the Eye Unit DES at the appropriate time. When this entity exits the clinic, the associated agent returns to the social care agent-based model, with their embedded sight progression model updated accordingly.

Up to this point, Joe had been coding the model himself although the conceptual modeling decisions were taken jointly by the CLC team members directly engaged on this project, namely the authors of this paper. However he found he had reached the limits of his expertise in Java, and so Stuart offered to help with the model coding. The primary issues were in cripplingly slow model performance, and the fact that AnyLogic encourages a visual development style (using building blocks customized via simple Java expressions) which focuses on ease of development and model visualization. This tends to result in a tightly-coupled code structure, where it is hard to separate out and abstract functionality without using ‘raw’ Java code. There are also occasions where the visual constructs either did not cover the required functionality (e.g., for the scheduling of appointments, which is essentially a classic algorithm with complex data structures), or were not quite the right fit conceptually (e.g., using statecharts to represent several related attributes changing over time), resulting in overly-complex and error-prone code. Joe’s background in DES and SD also naturally meant that he turned to DES constructs for some of the internal ‘plumbing’ (e.g., the arrivals of agents to the clinic), when there were actually better-suited AnyLogic mechanisms (collections of AnyLogic dynamic events for all a patient’s appointments in this case).

We thus altered the code to use AnyLogic constructs where they best fitted, supplemented with Java code where necessary. The performance issues turned out to be largely due to the use of multiple SD models. The difference equations which underlie the SD models were recoded in raw Java, resulting in performance improving by a factor of over 100. (A simulated year, starting with 500 agents and ending with approximately 650, changed from taking around 35 minutes to taking around 8 seconds for the core model, or around 12 seconds fully visualized with statistics-gathering.) To enable the loosely-coupled architecture which we wanted to swap-in and out components, we had to ‘subvert’ the standard AnyLogic structure somewhat to allow for a layered architecture where sub-components could interact via agreed interface definitions (rather than direct coupling). The layered structure also allowed for a clean separation of the core model from statistics gathering and visualization, which allowed us to streamline performance by only including aspects that were needed in each run. Hence we were able to perform both visualized and non-visualized single runs.

In addition to the above-mentioned face validity checks using the animated version of the clinic model, we have undertaken some further preliminary validation. The Eye Unit clinicians agreed that the pathways followed by different AMD patients (new, follow-up and injection) are correct. They also confirmed the probabilities used to route non-AMD patients to various procedures. The care need levels and provision levels are consistent with the data obtained from the ELSA study (2012). Future planned validation includes a statistical comparison of the model results in terms of the Eye Unit performance (waiting time distributions, process times, length of stay, number of patients seen, number of patients who miss appointments) against historical data which will be requested from the Eye Unit once ethical approval has been received. We shall also check the clinical progression of AMD against the literature used to populate the model to ensure that sufferers progress at the expected rate (within a given range). We shall check that

the social care and social care provision levels at an individual level (and therefore also at population level) are consistent with the data extracted from ELSA (2012) and other sources.

This study is still work in progress. To date, the model has mainly been used to conduct hypothetical experiments to investigate the relative effects of different improvement strategies. At the time of writing, we are in the process of acquiring cost information and are preparing to undertake a patient survey in the clinic to gather better data on how patients travel to the hospital. We are also talking to the relevant Local Authority in order to increase engagement from the social care sector, and to charities such as the Royal National Institute for Blind, who are very interested in this research.

7 REFLECTIONS

In the terminology of Chahal and Eldabi (2008), this is an *integrated* hybrid model. The different components of the model are intertwined within the same software implementation and (in theory at least!) there is seamless transition between the different sections and paradigms within the model. As such, it could be said to be a truer hybrid than the chlamydia model, which Chahal and Eldabi would term a *hierarchical* hybrid: the two components are distinct and one is a subset of the other. Arguably, it is easier to do this between DES and agent-based modeling than it is between DES and SD, where there are many technical problems in moving between an individual-based and a mass-based approach (Brailsford et al., 2010).

A number of modeling decisions were taken at various points: some driven by “client” need, some by particular aspects of the system being modeled, and some by the constraints imposed by the software tool we were using. This would be the case whether one were building a single-paradigm model or a hybrid, and would, in an ideal world, affect the choice of modeling approach: although we suspect that most people tend to stick to the approach (and software tools) they know best.

An expert programmer will almost certainly find commercial software packages frustrating. They force the user to think in a particular way and often conceal and obfuscate things that an expert user would like to be able to modify. Brailsford’s early DES models for HIV/AIDS and diabetic retinopathy were coded (almost) from scratch in Borland Delphi, using a library of routines for sampling, etc, and at the time, the early 1990’s, she would not have dreamed of using Simul8 or indeed any other commercial package. However there is always a trade-off between flexibility and convenience, and software has made many advances in the past twenty years!

As stated earlier, this paper is not intended to be a systematic critique of AnyLogic, but more of a narrative description of the process of building a hybrid model. The team reflect on their personal experience below.

Joe Viana said: *“Initially when I developed the model, I tried to use as many of the default functions as possible. This had obvious limitations, mainly how long it took to run. This was entirely due to the number of SD models which represented the agents’ eyes. I think given time, it may have been possible to build the model entirely in AnyLogic, but that would have required learning the AnyLogic way, of combining models, and we wanted to combine models in a particular way. I think we have broken away from the default functionality in almost every aspect apart from the DES models. I believe there is an expectation in AnyLogic that the agent part of the system should be modeled using statecharts, but this could have contributed towards the performance issues we encountered, resolved by Stuart’s recoding of the model. Even though we didn’t break away from the DES functionality I found that the most frustrating, quite possibly as I was so familiar with Simul8. The detachment of the visualization of the DES from the actual model frustrated me. It may have been easier to code the model from scratch, but that would have been beyond me. I think it may have been easier to combine an SD model (or series of SD models) built in VENSIM/Stella/etc, with a DES of the Eye Unit, built in Simul8/ARENA/etc, with an Agent based model built in Python/Java. I think it would also have been interesting and insightful to build the model entirely in one paradigm, either DES, Agent based or SD, as each would have given a different perspective on the system. The reason that we didn’t was discussed in the 2012 WSC paper (Viana et al., 2012): we used what we deemed to be the most appropriate tools for the job, based on subjective choice, modeling inter-*

est, and fit with CLC objectives. In fact I had already developed a causal loop diagram of the entire system and was planning to develop a simple SD model.”

Stuart Rossiter said: *“AnyLogic’s visual-based model construction and extensive visualization capabilities make it very useful for fast development of models, and in getting stakeholder buy-in via compelling presentations. However, there is a threshold of model complexity beyond which it is important to understand what is going on ‘under the covers’ of the software tool, and to be able to write one’s own extensions and alternatives as needed. In a hybrid model this is exacerbated because the combination of techniques (and related modeling perspectives) mean that the number of possible design decisions multiplies. No tool is likely to be able to foresee all of these, or design building blocks generic enough to be reusable in all circumstances. Thus, it is hard to avoid the need to become highly competent in the chosen tool and, where it is extended via ‘raw’ code (like in AnyLogic), this effectively means being a strong programmer in the normal software development sense.*

AnyLogic’s proprietary nature and tiered set of editions was also problematic for us, because only the most expensive edition (AnyLogic Professional) allows one to export AnyLogic models as reusable Java libraries, and thus interface them with code written using other tools. (Of course, one can still do this indirectly via file-based interfaces or similar.) This also hampers attempts to develop the model in a more principled software development manner: for example, since the model and all its experiments are included in a single AnyLogic-format file, multiple developers cannot easily develop features (or experiments) separately and merge the changes later.”

Amos Channon, as a demographer and social statistician, had a somewhat different perspective. He commented: *“I am extremely interested in the technicalities of the model mainly out of natural curiosity – but I am also very aware of my limitations in understanding what has actually been done. The ideas that lie behind the model are very difficult to apply to my usual research – these problems don’t seem to be addressed in demography. So at the moment it seems like a black box: you input the data and relationships and out pop the results. I can’t comment on the workings of the box, although I would love to be able to.*

It has most definitely changed my views of simulation. I think that it has made me far less scared of it – it is the usual story of garbage in garbage out – it doesn’t matter how much effort you put into the model structure, but if you start off with poor quality data and theory then the model won’t tell you anything. What I do really like is the way to link different systems together, which is something that basic social statistics can’t do. It is also the dynamic nature of things rather than static statistical models. I am interested in how much you have to feed the engine of this model: there are a lot of parameters in the model that we should specify better based on real data. This takes a lot of time to do properly, and I am unsure of how much impact it will have on the final outcome, but this is really interesting to me.”

Andrew Lotery, as a clinician, commented *“It has been interesting to see the model evolve and I hope in the long term it will help us improve clinic flow and our patients experience. We are constantly changing our practice by developing new clinic space and techniques to deliver treatment. It will be interesting to test the model against our real life experience. We see over 70000 people a year in our clinics so critical analysis of our work patterns is essential to cope with this ever increasing demand.”*

In conclusion we must reflect on whether the hybrid simulation modeling paradigm chosen to study the links between health and social care for AMD was a good choice, or if it has indeed been “more trouble than it’s worth”. The answer to this differs depending on the hopes and expectations of the individuals involved, as well as the level at which this question is examined.

From the perspective of those involved with no experience of non-hybrid models, it is difficult to appreciate if hybrid simulation is more trouble than it’s worth – there is nothing to compare it to. With interdisciplinary collaborations becoming more common it is clear that this experience will not be unique, with researchers using new methods such as SD, DES and agent-based modeling for the first time. These researchers will have to rely on experienced operations researchers and other modelers to guide the choice

of paradigm, based on what is considered to be most applicable to the problem area. In this case the research questions posed highlighted the links between the health and social care systems, so a hybrid model is appropriate. As a result the hybrid model is the “way forward”. However, from the perspective of staff at the Eye Unit it is arguable that the more complex hybrid model is adding a layer of complexity that is not needed nor desired when the initial brief was to consider a better configuration of the department. A simple DES model would have been applicable – and have taken less time and effort to construct. Although the wider social environment is appreciated as vital by all involved in the study, the addition of this to the model and the hybrid simulation has delayed any recommendations for configuration changes.

In terms of model design, the high-level combination of paradigms we used seemed an intuitively sensible “best tool for the job” fit to the subsystems that they covered (clinic, ‘socially-embedded’ individuals, and sight loss over time). The interpretation of “best tools” is always subjective, but what is important is that the assumptions made are generally accepted by the modelers and stakeholders. Validating outcomes against empirical data is always possible, but we are often dealing with what-if scenarios which makes this less relevant. Nothing to date has changed our mind on the relevance and value of the hybrid architecture.

Having said that, the ultimate methodological goal is to include elements which best leverage the core concepts of the individual paradigms, whilst also showing interesting ways to combine them at different levels of granularity; always with the strong caveat that these elements add insight into the health and social care domain (and are not being done as a clever academic exercise). As it stands, the current model still has some way to go here, but it provides a useful platform from which to do so, and each of us has different interests in this area. For example, we could: (a) explore interaction-driven social aspects, where care relationships form endogeneously via explicit decisions on migration, care-giving, care-needs-perception, etc.; (b) enrich the sight model with links to the social activities of the agent (either direct influences, such as smoking, or in the attitudes to sight loss); and (c) add spatiality, with links to the spatial distribution of social attributes (e.g., poverty or social norms). As we said earlier, we would also expect to compare alternative models: we are not adding detail to fashion an increasingly accurate or plausible single model. This includes comparing with single-paradigm alternatives to the entire model.

In terms of model implementation, it is difficult to separate our backgrounds from our judgments, so each of us has had different reactions to the use of AnyLogic, and perceive different strengths and weaknesses. Combining paradigms in a single package is clearly useful, and avoids the need for extra “plumbing” code. However, the proprietary nature of AnyLogic worked somewhat against us in creating a reusable, modular model. To some extent, this would apply to any commercial tool. The current model could perhaps have been coded with less overall effort in one of the single-paradigm packages that one or more of us are proficient in, even given our subsequent experience with AnyLogic. However, some of the things we intend to add, such as the examples above, should make the strengths of using differing paradigms more apparent, and make the coding more intuitive and concise in AnyLogic (compared to using a single package). That is, future model structures will increasingly leverage the expressive power of the hybrid approach.

ACKNOWLEDGMENTS

We gratefully acknowledge the many helpful discussions with the remaining members of the Care Life Cycle research team not named explicitly as authors on the paper: Seth Bullock, Jane Falkingham, Maria Evandrou, Jonathan Klein, Jason Noble, Jakub Bijak, James Robards, Eric Silverman and Athina Vlachantoni. We also gratefully acknowledge support from the EPSRC under grant EP/H021698/1. Finally, we would like to express our gratitude to the staff of the Eye Unit at University Hospital Southampton.

REFERENCES

- Access Economics PTY Limited. 2009. Future sight loss UK (1): The economic impact of partial sight and blindness in the UK adult population. RNIB.
- Age-Related Eye Disease Study Research, G. 2000. Risk factors associated with age-related macular degeneration- A case-control study in the age-related eye disease study: age-related eye disease study report number 3. *Ophthalmology* 107: 2224-2232.
- Age-Related Eye Disease Study Research, G. 2005. A Simplified Severity Scale for Age-Related Macular Degeneration. *Archives of Ophthalmology* 123: 1570-1574.
- AnyLogic, www.xjtek.com/AnyLogic, accessed 07.09.13
- Brailsford, S.C., M.S. Desai and J. Viana. 2010. Towards the holy grail: combining system dynamics and discrete-event simulation in healthcare. *Proceedings of the 2010 Winter Simulation Conference, Baltimore, Md.* B. Johansson, S. Jain, J. Montoya-Torres, J. Hagan, and E. Yücesan, eds.
- Brailsford, S.C., Silverman, E., Rossiter, S., Bijak, J., Shaw, R., Viana J., Noble, J. Efstathiou S. and Vlachantoni A. 2011. Complex systems modeling for supply and demand in health and social care. *Proceedings of the 2011 Winter Simulation Conference, Phoenix, Az.* S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, and M. Fu, eds.
- Brailsford, S.C., Evandrou M., Luff R., Shaw R., Viana J., Vlachantoni A. and Willis R.M. 2012. Using system dynamics to model the social care system: simulation modeling as the catalyst in linking demography to care delivery. *Proceedings of the 2012 Winter Simulation Conference, Berlin, Germany.* C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A.M. Uhrmacher, eds.
- Brailsford S.C., Harper P.R., Patel B and Pitt M. 2009. An Analysis of the Academic Literature on Simulation and Modelling in Healthcare. *The Journal of Simulation*; 3:130-140.
- Care Life Cycle. 2013. "The Care Life Cycle [Project]." Accessed 24 March, 2013, from <http://www.southampton.ac.uk/clc>.
- Chahal, K. and Eldabi, T. 2008. Applicability of hybrid simulation to different modes of governance in UK healthcare. *Proceedings of the 2008 Winter Simulation Conference, Miami, Florida.* S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, and J. W. Fowler, eds.
- Courgeau, D. 2007. Multilevel synthesis : from the group to the individual. Springer.
- Department of Health, 2003. http://webarchive.nationalarchives.gov.uk/+/www.dh.gov.uk/en/publicationsandstatistics/publications/publicationspolicyandguidance/dh_4009653, accessed 11.05.13.
- ELSA. 2012. English Longitudinal Study of Ageing. Accessed 1st June, 2012, from <http://www.ifs.org.uk/ELSA>
- Gilbert, N. and K. G. Troitzsch, 2005. Simulation for the Social Scientist (2nd ed.). Open University Press.
- Jun, J. B., S. H. Jacobson, J. R. Swisher, 1999. Application of discrete-event simulation in health care clinics: A survey. *Journal of the Operational Research Society* 50: 109-123.
- Morecroft, J. and Robinson, S. 2006. Comparing discrete event simulation and system dynamics: modelling a fishery. *Proceedings of the 2006 OR Society Simulation Workshop SW06*, 137-148.
- Noble, J., E. Silverman, J. Bijak, S. Rossiter, M. Evandrou, S. Bullock, A. Vlachantoni and J. Falkingham. 2012. Linked lives: The utility of an agent-based approach to modelling partnership and household formation in the context of social care. *Proceedings of the 2012 Winter Simulation Conference, Berlin, Germany.* C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose and A. M. Uhrmacher, eds.
- Tako A.A. and Robinson S. 2009. Comparing discrete-event simulation and system dynamics: users' perceptions. *Journal of the Operational Research Society*, 60:296-312.
- Viana, J., Rossiter, S., Channon, A.R., Brailsford, S.C and Lotery, A.J. 2012. A multi-paradigm, whole system view of health and social care for age-related macular degeneration. *Proceedings of the 2012*

Winter Simulation Conference, Berlin, Germany. C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A.M. Uhrmacher, eds.

Viana, J. 2011. The development of a combined simulation approach in a sexual health context: combining discrete event and system dynamics simulation to form a composite model, PhD thesis. University of Southampton, UK.

Wooldridge, M. J. 2002. *An introduction to multiagent systems*. John Wiley & Sons.

AUTHOR BIOGRAPHIES

SALLY C. BRAILSFORD is Professor of Management Science at the University of Southampton, UK. She received a BSc in Mathematics from the University of London, and MSc and PhD in Operational Research from the University of Southampton. Her research interests include simulation modeling methodologies, system dynamics, health service research and disease modeling, and the modeling of human behavior in healthcare systems. She is chair of the European Working Group on OR Applied to Health Services (ORAHS) and is an Editor-in-Chief of the journal Health Systems. She is on the editorial boards of Health Care Management Science, the Journal of Simulation, and Operations Research for Health Care. Her email address is s.c.brailsford@soton.ac.uk.

JOE VIANA is a Research Fellow in the Southampton Management School at the University of Southampton. He received a BSc with Honors in Sport and Health Science with Psychology from the University of Southampton, and MSc and PhD in Management Sciences from the University of Southampton. He is primarily interested in modeling health care systems, and in combining different simulation paradigms at different levels. His email address is J.Viana@soton.ac.uk

STUART ROSSITER is a Research Fellow on the Care Life Cycle project at the University of Southampton. Originally a mathematician, he worked as a software developer and architect for 10 years before focusing on social simulation via an MSc in Multidisciplinary Informatics (Leeds University), and a PhD in the simulation of electricity markets (Strathclyde University, ongoing). He is interested in general issues of modeling and simulation methodology, particularly in processes, software frameworks and agent based modeling. His email address is S.Rossiter@soton.ac.uk.

ANDREW ‘AMOS’ CHANNON is a Lecturer in Demography in the Division of Social Statistics and Demography at the University of Southampton. After a BSc at the University of Durham in Psychology he was an IT Project Manager for 6 years before completing an MSc and PhD in Social Statistics at the University of Southampton. His main research interests are in child health in sub-Saharan Africa and population health in the Middle East, but also researches in health systems around the world, especially related to the supply of health care. His email address is a.r.channon@soton.ac.uk.

ANDREW LOTERY is Professor of Ophthalmology in the Faculty of Medicine at the University of Southampton. He runs an integrated laboratory and clinical ophthalmic research program, dividing his week between directing vision research in the Gift of Sight research laboratories at The University of Southampton and his clinical responsibilities seeing patients in Southampton Eye Unit. His special interests are in medical retina with major interests in age related macular degeneration and ophthalmic genetics; stem cell & synthetic biology repair, and gene therapy. As well as being Editor in Chief of the journal Eye he and his team have published in high impact journals such as Nature Genetics, the Lancet and the New England Journal of Medicine. His email address is a.j.lotery@soton.ac.uk.