Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

COMBINING BOOTSTRAP-BASED STROKE INCIDENCE MODELS WITH DISCRETE EVENT MODELING OF TRAVEL-TIME AND STROKE TREATMENT: NON-NORMAL INPUT AND NON-LINEAR OUTPUT

Kim Rand-Hendriksen
Joe Viana
Fredrik Dahl

Health Service Research Centre,
Akershus University Hospital
Sykehusveien 25,
1478 Lørenskog, NORWAY

ABSTRACT

Incidence rates in simulation models are often assumed to stem from Poisson processes, with rates based on analyses of real-life data. In cases where the record of data is limited, or observed rates are low, the stochastic process involved in sampling from modeled distributions may not adequately reflect the uncertainty around the estimated input parameters. We present a conceptually simple, but computationally demanding, method for generating variance in incidence through the use of bootstrapping; for each subsample, a regression model is fitted, and the simulation model is run repeatedly sampling from the fitted model. Stochasticity is introduced at two levels; data for fitting the regression, and sampling from the fitted model. We illustrate this hybrid approach using Norwegian stroke records to generate stroke incidences with age, sex, and location, in a simulation model made to analyze travel time, queuing, and time to treatment in regional stroke units.

1 INTRODUCTION

It is well known that in situations involving queuing, such as waiting lines at shops, or new arrivals at hospitals, the waiting time can change substantially when the demand reaches certain thresholds or tipping points (e.g. Davis, Eisenhardt, and Bingham 2007). Such non-linear responses are often difficult to represent and handle using regular statistical methods. One of the main comparative advantages of simulation models is that they can be made to mimic real-life queuing response to resource depletion, and such tipping points emerge. In cases where queuing systems behave very differently below and above certain thresholds in demand for resources, accurate representation of variation in demand is important to model validity; given an average demand below the tipping point, insufficient variance will likely result in the tipping point being reached less often than in reality, and exaggerated variance will result in the tipping point being reached too often. If the intention of the model is to inform decision making regarding resource expenditure for systems where the cost of queuing is high, appropriate handling of variance in demand is crucial.

In cases where it is reasonable to assume that incidences are independent or close to independent, time between events is often modeled as a Poisson point process. Assuming that events can be modeled as stemming from a Poisson process is attractive, because it only requires a single parameter, reflecting the rate or intensity per unit of time, to perform stochastic sampling of time between events. Since the average rate of events is often relatively easy to measure or estimate, the Poisson process is regularly applied in a wide variety of disciplines concerning random phenomena, including simulation modeling. Consequently, methods and applications involving Poisson processes have been the topic of several papers at previous Winter Simulation conferences, with several hundred manuscripts using the term *Poisson* at least once (approximately 900 listed in Google Scholar, and 162 unique PDF files under the papers sections in the winter WinterSim archives (Google Scholar 2017, Google 2017). Importantly, when average incidence rates are estimated, the tendency is to use the best fit point estimate(s) as the sampling parameter(s) when performing random sampling in simulation models, with varied figures sometimes being introduced as a form of sensitivity analysis. Given the tendency for queuing systems to display non-linear response to linear changes in demand or frequency, even modest misestimation and sampling error for incidence rates can, in certain cases, have substantial impact on overall model behavior. Henderson (2003) exemplifies this, and discusses various methods for handling input uncertainty. Morgan et al. (2016) provide a short summary of the literature on input uncertainty, and discuss methods for analyzing and quantifying the uncertainty in question. The procedure discussed in this paper is an implementation of a resampling method used to reflect input uncertainty in a hybrid model setting.

In this paper, we discuss methods for representing variance in the number and frequency of strokes in Norway, in an early and simplified version of a simulation model that will eventually be used to model different scenarios for the placement, staffing, and management of stroke units in the country. The modeling of strokes and stroke units represent a particularly interesting challenge. First, it is well established that time is of the essence for the clinical outcome for patients with stroke, as the region of the brain suffering irreversible harm grows larger on a time scale of minutes. As clinicians describe it, 'time is brain' (Kamal et al. 2014; Saver 2006). The clinical literature strongly indicates that the clinical outcome is improved in patients admitted to larger, specialized stroke units than less specialized units, or hospitals in which the staff has less experience with stroke. Since strokes are relatively rare, this creates an interesting optimization problem, particularly in sparsely populated regions, since the need for experienced clinicians in larger clinics conflicts with the need for short travel time to the place of treatment. In addition to these considerations comes challenges brought on by resource constraints; the equipment required to perform state-of-the-art treatment is expensive, and maintaining a constant capacity in terms of staffing, hospital beds, and equipment is costly. It is challenging to determine the ideal capacity of stroke units, since the number of patients can vary substantially, and maintaining a capacity that is never exhausted may be prohibitively expensive.

We combine two main modeling components: a regression-based statistical model for predicting the rate of strokes by age, sex, and region, implemented in the R statistical package (R Core Team 2016); and a discrete-event/agent-based hybrid simulation model implemented in AnyLogic (2017), in which stroke patients agents are spawned in each municipality based on random samples from an exponential function from the regression model, to geographic locations on a map, from which they travel along the road network to their respective stroke units. These models will be described in greater detail in the methods section. However, the important component for this paper lies in the first model, in which the rate of strokes is estimated based on patient records for all stroke patients in Norway between 2010 and 2015.

When parameters of simulation models, such as the rates of stroke in this model, are based on statistical analysis of observation sampling or historic records, the result will generally be a set of fixed parameters with estimated standard errors. Unless otherwise specified, the convention in estimation of standard error estimates is done under assumptions of normality and homoscedasticity. When normality cannot be assumed, or deviation from normality is to be expected, a variety of methods have been developed to enable estimates of uncertainty, including bootstrapping. In its simplest form, a bootstrap analysis of a sample with size N consists of a sampling a large number of subsamples, each with the same size as the original, with replacement, from the initial sample. A bootstrap can be construed as a simulation of variance, generated through resampling from the observed sample. Compared to regular statistical estimation of variance based on the assumption of normality, bootstrapping is computationally demanding. However, the procedure is highly versatile, since it does not require assuming any specific underlying distribution for the variation in the data. In this paper, we present a demonstration case in which a bootstrap procedure is used prior to the regression-based estimation of mean incidence rates for strokes. In the following, we will first present the simulation model used to represent stroke patients, travel to stroke units, and treatment in stroke units. Second, we present the model used to generate estimates of mean stroke incidence rates. Third, we describe the interface between the incidence model and the simulation model, and two scenarios for comparison. We then present results, before moving on to a discussion of the strengths and weaknesses of the described modeling approach.

2 STROKE PATIENT SIMULATION MODEL

2.1 Overview

The stroke patient simulation model is a DES/ABM hybrid under development (see Figure 1 for a conceptual model), intended to be used to test various scenarios for future placement and organization of stroke units in Norway in response to forecasted changes in settlement distribution and demographic makeup of the population. The model is implemented in AnyLogic, and contains several nested models for various sub-elements, the most important of which are the patient agents, hospitals, stroke units, and the geographic model including municipalities. The general approach used to combine discrete-event and agent-based (and system dynamics, if relevant) modeling properties has been described elsewhere (Viana et al. 2016). The model handles patients as entities with assigned properties, which is common in agent-based modeling. However, the patients have no real agency, and are passively moved around in a structure-based geographical model and queues in the hospital, which is more typical of DES. Whether the model would be best categorized as DES or as ABM is not clear, and we do not find this categorization to be particularly helpful. We consider our approach to modeling to be eclectic, with characteristics drawn from various classical approaches to simulation modeling.

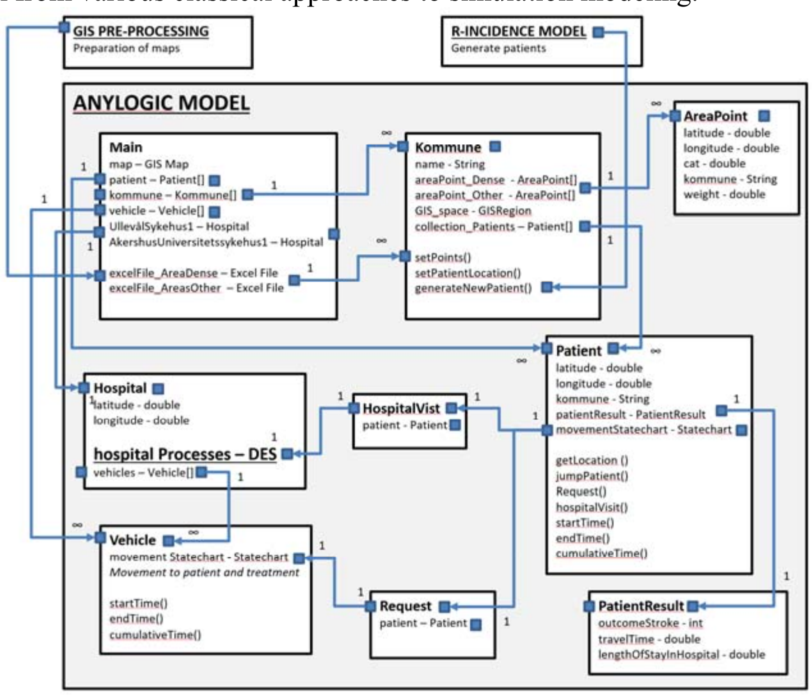

Figure 1: Conceptual model.

2.2 Patients

Patients are represented as agents with a large number of properties, including age, sex, place of residence, and stroke type/severity classification, and other diseases (Figure 2 shows the patient agent represented in AnyLogic). As we are in the early stages of development, and the current object of interest is related to the representation of stroke incidence, we have simplified patient type to vary by type of stroke plus geographic location only, disregarding age and sex. When spawned, patient entities are randomized to the following stroke or stroke-like categories: transient ischemic attack (TIA, a very small stroke with limited long-term consequences, approximately 17% of patients), obstructive stroke (cerebral embolism, about 42.5%), hemorrhagic (7.5%), and "stroke mimic" (33%). These proportions reflect observations from an ongoing study of all patients admitted with suspected stroke at Akershus University Hospital, Norway. The stroke incidence model (see section 3) provided rates for obstructive and hemorrhagic strokes only, and patients with TIA and stroke mimics were spawned separately, based on mean rates that did not vary by scenario. Stroke mimics are patients who are admitted with stroke-like symptoms, but who turn out not to have stroke. These typically spend less time in treatment than actual stroke patients, and have good clinical outcomes. However, since they contribute to queuing and occupy costly treatment resources in the investigative phase of their hospitalization, their inclusion is important. Patients agents are randomly assigned a starting location within a densely habited zone or a building within their municipality, from which they travel by the fastest available route to the stroke unit for which the catchment area covered the relevant municipality.

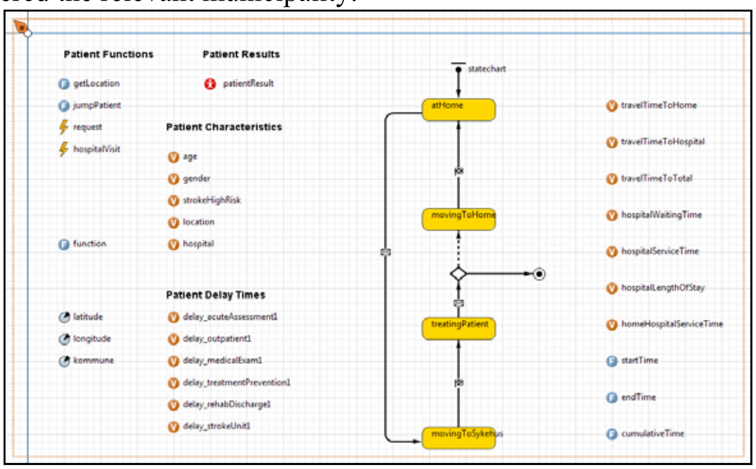

Figure 2: Patient agent. The DES/ABM hybrid places patient-centered characteristics in the agent, while some routing logic is nested in structure such as hospitals and municipalities.

2.3 Hospitals and stroke units

Hospitals and stroke units are defined based on a generic platform in which it is possible to specify the available wards and units, and to tailor the treatment paths, available procedures, number of hospital beds, etc. For this specific study, we restricted modeling to two hospitals: Ullevål University Hospital, covering Oslo, the capital of Norway, and Akershus University Hospital, with a catchment area to the north, east, and south of Oslo (see figure 3 for the placement of the two units, and the outline of the combined catchment area). These two units have a joint catchment area of roughly 1.1 million inhabitants, and treat around 2500 new strokes each year, split at roughly 1500 for Ullevål and 1000 for Akershus. In this study, the stroke units and treatment paths within these are purposely kept simple: Ullevål was assigned a total of 38 beds, and Akershus 28. In reality, the number of available beds varies, but 38 and 28 are reasonable approximations. When admitted, length of stay (LOS) is set to a minimum of 2 hours, after which the subsequent stay is drawn from an exponential distribution, with the mean depending on the type of patient: 2 days for stroke mimics, 3 days for TIA, 5 days for obstructive, and 7 days for hemorrhagic strokes. If all beds are occupied, new arrivals would be set to queue.

2.4 Geography, location, and travel to stroke units

The geographic model component is a GIS-based representation of Norway based on very extensive map data publicly available at the geoportal "GeoNorge", www.geonorge.no (Giff et al. 2008). From files provided by geonorge, we extracted information about road networks, municipality borders, and habitation in the form of densely populated areas and larger buildings. Based on mean stroke rates for each municipality, derived from the incidence model (details in section 3), stroke patients were randomly spawned within habitable zones, from which they would start moving to their respective stroke unit. The model supports simulation of each patient finding his/her way to the closest road, after which the agent would follow the road network at the speed limit to the hospital. The model also supports the generation of ambulances that would use the road network to pick up the patient and transport him/her to the stroke unit. Unfortunately, simulations of following the road network are computationally demanding in the extreme (each 1 year run requiring about 200 seconds of CPU time), and we required a very large number of runs (approximately 45 000) for our analyses. As the subject of this test case analysis is more dependent on variation in travel time than on fully accurate representation of travel time for each agent, we simplified the model to transport agents to the stroke unit following a straight line, which reduced model run time by a factor of about 50.

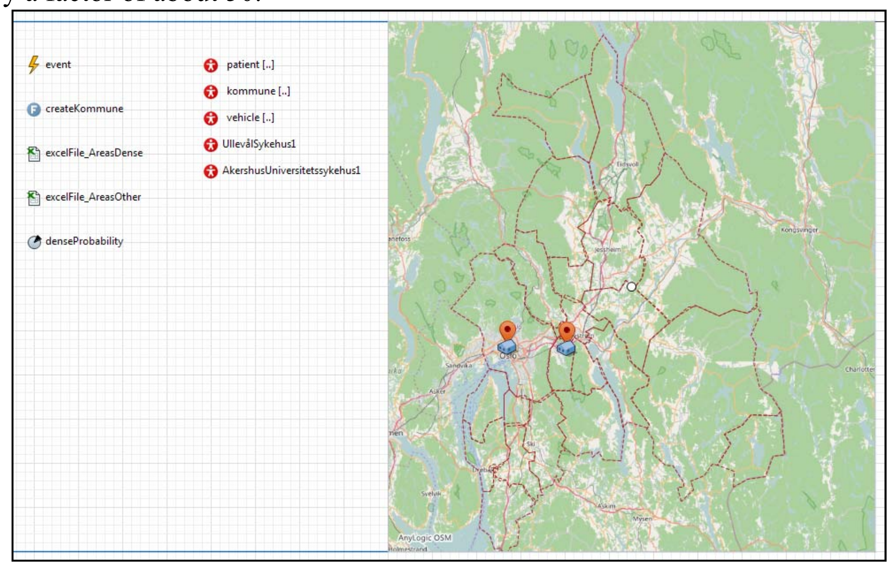

Figure 3: The catchment area for Ullevål (to the left) covers only Oslo, while the rest of the areas inside the dotted lines describe the municipalities in the catchment area for Akershus University Hospital.

3 STROKE INCIDENCE MODEL

3.1 Main regression model

To predict stroke incidence, we used anonymized data from the Norwegian Patient Registry and the Norwegian Cause of Death Registry, in the form of a combined dataset with all hospital records for 2010-2015 for every person in Norway with one or more recorded stroke diagnoses in either registry within that time frame. Strokes were defined by ICD-10 codes G45 (excluding subcode 4), I61, I63, I64, I69 (excluding subcodes 0 and 2), and H34.1. After reducing the data to a single record per unique stroke, we were left with a total of 100645 recorded strokes, combined with information about time, age, sex, comorbidities, and municipality/city district. Through an API provided by Statistics Norway, we extracted information about the demographic makeup of the Norwegian population in terms of yearly age and sex for each municipality/city district (tables 07459 and 10826) for each year from 2010-2017 (Statistics Norway 2017a, 2017b). From the stroke records and the tables from Statistics Norway, we determined the total population and corresponding number of strokes by year (year-2010), municipality, sex (dummy for female), and age bracket (dummies for ages 0-5, 6-20, 21-30, 31-40, 41-45, 46-50, 51-55, 56-60, 61-65, 66-70, 71-75, 76-80, 81-85, 86-90, 90-95, and 96→), for a total of 192 groups. With this data, we fitted a relatively simple Poisson regression model:

$$\log(\lambda_i) = \beta_{female} * female_i + \beta_{age} * age_i + \beta_{year} * year_i + \log(population_i)$$

Thus, λ is the estimated yearly stroke rate per individual, broken down by age, sex, and year. The model was fitted in R using the built-in *glm* function with a Poisson family:

```
glm(strokes ~ -1 + female + age + year + offset(log(pop)),family=poisson, data=df)
```

The regression was fitted without an intercept (the initial -1 in the R call), as each age bracket is assigned an independent value. Using this model, we predicted stroke rates for each municipality/city section in Norway for 2017, based on the recorded population by yearly age and sex for January 1st 2017 provided by Statistics Norway. The predicted age and sex-specific rates were aggregated by municipality/city section, to be used for generating stroke cases in AnyLogic. This model produces mean rates for obstructive and hemorrhagic strokes only. Stroke mimics and patients with TIA were spawned sampling from mean rates that did not vary between scenarios.

3.2 Handling measurement error with bootstrapping

Stroke rates are relatively low, at roughly 100 000 strokes over 30 million person years. Stroke risk increases rapidly with advancing age, but the demographic groups at risk diminish rapidly in size due to overall age-related mortality. Consequently, observed stroke rates by regions such as municipalities are highly erratic, and it is difficult to determine how long an observation period, and how large a population, are needed to ensure a stable set of observations. The Poisson regression model provides information about the standard error of measurement (SE) for each parameter, and it is possible to generate SE-estimates for predicted log-rates. One way in which this estimated uncertainty could be taken into account in a simulation model would be to use the SE estimates to conduct random sampling around the estimated stroke rate for each municipality * age * sex group, prior to random sampling from an exponential distribution for the group. As this procedure adds a new layer of stochasticity to the model, the number of runs required to produce stable and reliable results is increased considerably. However, this procedure rests on the assumption that the (log) variance in stroke rates is adequately described using a normal distribution, which is unknown. Since we do not know what distribution to expect, we use bootstrapping.

As described under section 3.1, we have exact information about the composition of the Norwegian population by year, municipality/city sector, sex, and yearly age; as well as the corresponding information for each recorded stroke. This allows us to generate a dataset matching every individual in the recorded Norwegian population (roughly 5 million) for the 6 years in question (i.e ca. 30 million observations), each with associated age, sex, municipality, year, and whether he/she had a stroke that year. With this dataset, we created 1000 group-based bootstrap subsamples. We defined subgroups of the population by age bracket (0-5, 6-20, 21-30, 31-40, 41-45, 46-50, 51-55, 56-60, 61-65, 66-70, 71-75, 76-80, 81-85, 86-90, 90-95, and 96→), sex, year, and municipality. Within these subgroups, we performed resampling with replacement, each resample of the same size as the group in question. This procedure introduces random variation in terms of stroke rates for each group, but does not alter the recorded composition of the modeled population in terms of age or sex within the municipalities. For each of the 1000 subsamples, we fitted the regression model described in section 3.1, and predicted stroke-rates for each municipality/city sector in 2017.

4 SCENARIOS AND INTERFACE

4.1 Scenarios

We were concerned with the potential impact of variation and uncertainty in the data used to estimate the stroke rates used to populate the stroke simulation model. As described in section 3.2, we decided to use bootstrapping to model variation in the records of historic stroke incidences. As a test of impact, we compared the distribution of ward occupancy, queuing, and time to treatment from (A) 1000 runs of the simulation model in which stroke incidence was sampled based on predicted rates provided by the main regression-based incidence model, to (B) 42 runs for each of the 1000 bootstrap-based incidence models, for a total of 42 000 runs. The models were run on several computers over a period of roughly three days.

4.2 Interface

The procedure described in this paper relies on modeling at several levels. First, regular statistical modeling is used to generate a main stroke incidence model. Second, bootstrapping is used to simulate natural variation in the incidence of stroke prior to regression modeling. The estimated stroke rates (both the main model and the 1000 bootstrap-based models) were used to inform random sampling of incidence rates in the simulation model. While it might be technically feasible to perform all these steps in a single piece of software, AnyLogic is not well-suited to conducting the regression and bootstrap analyses, and simulation modeling in R is less streamlined. To draw on the strengths of each, we are developing an interface between AnyLogic and R. R is open source, and intentionally built to support direct interfacing with other software, while AnyLogic is proprietary, and not built with this in mind. Fortunately, AnyLogic is built on Java, in which we have written code to control R. As of yet, the interface is rudimentary, and requires substantial user input to work properly, but the intention is to continue development until it is possible to have AnyLogic to seamlessly request analyses from R while running simulation models, and to incorporate the results without user interference.

5 COMPARISONS AND RESULTS

Mean travel time and mean stroke rates did not deviate noticeably between scenarios A (incidence from the main regression model) and B (1000 bootstrap-based incidence models). However, there was slightly more variation in stroke rates in scenario B, resulting in a pattern of very slightly elevated peaks in queuing, and slightly increased mean time to treatment.

To compare waiting time between scenarios, we first determined the percentile of waiting times for the patients in each run (1000 for scenario A, 42 000 for scenario B). Thus, for each specific run, we could tell the waiting time for a patient in any specific percentile, e.g. the 50 percentile or 75 percentile. Second, in each scenario, we determined the percentiles across runs for each percentile of waiting time within runs. This allowed us to display the variation in waiting time graphically (Figure 4).

To understand the graphs, each line represents a percentile across runs, and the x-axis indicates the percentile within runs. If you consider the top line in the graph, it indicates the maximum observed waiting time for patients at specific percentiles within their respective runs. It crosses 50 and 75 on the x-axis at 0 and around 30, respectively, which tells us that the 50th percentile of waiting time within all runs was 0 (since the maximum observation was 0), and that the maximum observed waiting time for the 75th percentile in any run was around 30 minutes. The second line indicates the 99th percentile across runs, and the lowest line indicates the minimum observation, which indicates that for the runs with the least queuing, patients in the 90th percentile for waiting time waited for 0 minutes.

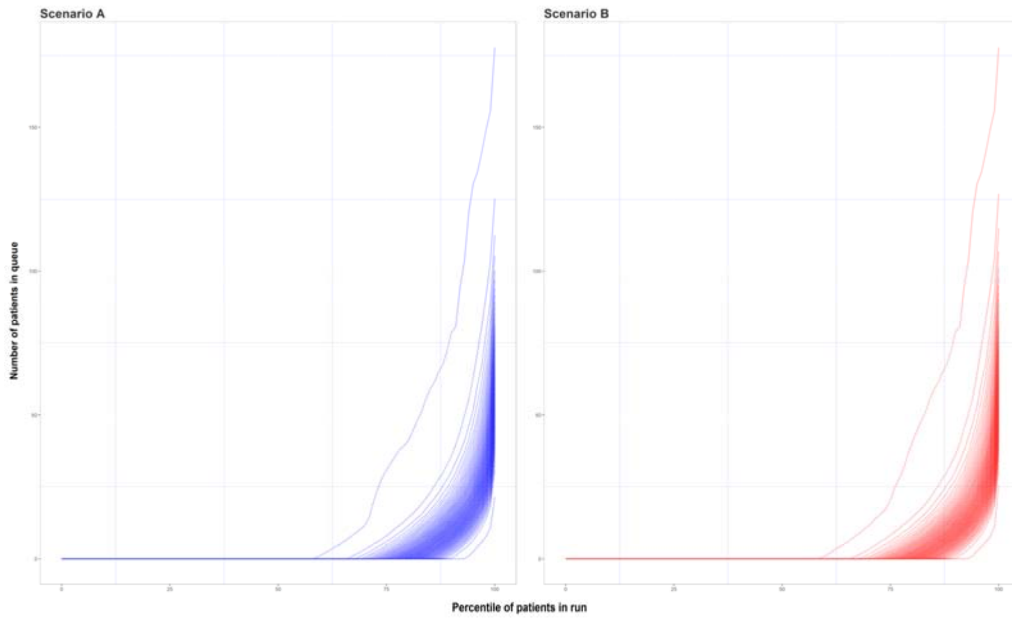

Figure 4: X axis is percentiles within each run. Y axis is number of patients in queue. Lines for each percentile across runs in scenario.

Queuing and occupancy were analyzed by stroke unit, and while both displayed a tendency to be higher for all measures in scenario B, the differences were tiny compared to the variation between individual runs.

6 DISCUSSION

Introducing variance in estimated mean stroke rates through bootstrapping resulted in marginally greater variability in stroke incidence, with slight impact on peak queuing, and peak time to treatment. However, the differences between the tested scenarios were all small. While the bootstrap procedure increased variance in stroke rates slightly, particularly when considering municipalities with small populations, the added variance was barely observable at the level of the full catchment areas for the two modeled stroke units. Considering that we operated with stroke as a binary variable, and the resource expenditure (time in treatment) per stroke was sampled independently, this should not come as a surprise; with this setup, the resampling could not produce clusters of outliers, able to occupy substantial resources and produce dramatic queuing. We observe the presence of a tipping point at which queuing becomes a problem, but the setup of the sampling in scenario B was not suited to exacerbate queuing substantially. Whether we would observe greater differences if stroke severity was included in the incidence prediction model, and resource expenditure was tied to severity, remains to be seen.

As reflected in the difference in the number of runs for the two scenarios, adding a new layer of stochastic sampling comes at a distinct cost in terms of computational efficiency, and this cost may be prohibitively high for many realistic implementations. In the scenarios tested in this study, the extra effort hardly seems worthwhile. However, a single finding of this sort is insufficient to conclude that disregarding potential input uncertainty is generally safe. It is likely that the volume of data per parameter in our prediction model was such that the variance introduced by the bootstrapping procedure all but drowned in the variance introduced by random sampling within each municipality. For modeling problems with less extensive empirical records, the proportional variance introduced by bootstrapping is likely to be larger.

The presented method may be too costly in terms of computation to be put to general use as a form of sensitivity analysis. However, there may be ways in which the required number of runs could be reduced to a more manageable level. For instance, depending on the form of the regression model (or other statistical modeling technique), it might be possible to rank the bootstrap-based models along some dimension of magnitude, and use only a strategic subset for sampling in the subsequent simulation modeling. For instance, if 1000 or more bootstrap models were generated, and the models are such that they allow for ranking, limiting the tested runs models falling on percentiles, or even deciles, would greatly reduce the computation time, while introducing most of the same variation. This relatively limited sampling might be sufficient to determine if handling input uncertainty through resampling is called for.

The models used in this study were simplified at many levels. Several of these simplifications are clearly unrealistic; for instance, in real cases of full occupancy, new stroke patients would either be diverted to other hospitals, or beds would be assigned from other wards, reserve personnel would be called in, and the ward in question would provide treatment at levels over optimum capacity. This is only one of many deviations from reality that need to be addressed before this particular stroke model can meaningfully inform decision making regarding stroke treatment in Norway.

7 CONCLUSION

We describe a procedure in which bootstrapping is used to handle input uncertainty in the estimation of stroke rates prior to simulation modeling. While it appears that the added level of stochasticity does result in altered overall model behavior in terms of queuing and waiting times, the impact was practically negligible in this study. The procedure used here is very computationally demanding, and might not be worthwhile for many applications, but there may be ways to reduce the computational burden to a more acceptable level.

REFERENCES

Davis, J. P., M. E. Kathleen, and C. B. Bingham. 2007. "Developing Theory through Simulation Methods." *Academy of Management Review* 32 (2): 480–499.

Giff, G., B. Van Loenen, J. Crompvoets, and J. Zevenbergen. 2008. "Geoportals in Selected European States: A Non-Technical Comparative Analysis." In *Conference, Small Island Perspectives on Global Challenges: The Role of Spatial Data in Supporting a Sustainable Future Location St. Augustine, Trinidad*, 25–29.

Google. 2017. "Google search for Poisson in WinterSim Archives." https://www.google.com/search?q=poisson+inurl%3A*papers+site%3Ainforms-sim.org+filetype%3Apdf

Google Scholar. 2017. "Google Scholar search for Poisson in WinterSim proceedings." https://scholar.google.com/scholar?as_vis=1&q=poisson&as_publication=winter+simulation&as_sdt=1,5

Henderson, S. G. 2003. "Input Model Uncertainty: Why Do We Care and What Should We Do About It?" In *Proceedings of the 2003 Winter Simulation Conference*, edited by S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, 90–100. Piscataway, New Jersey: Institute of Electronic and Electrical Engineers, Inc.

Kamal, N., O. Benavente, K. Boyle, B. Buck, K. Butcher, L. K. Casaubon, R. Côté, A. M. Demchuk, Y. Deschaintre, D. Dowlatshahi, G. J. Gubitz, G. Hunter, T. Jeerakathil, A. Jin, E. Lang, S. Lanthier, P. Lindsay, N. Newcommon, J. Mandzia, C. M. Norris, W. Oczkowski, C. Odier, S. Phillips, A. Y. Poppe, G. Saposnik, D. Selchen, A. Shuaib, F. Silver, E. E. Smith, G. Stotts, M. Suddes, R. H. Swartz, P. Teal, T. Watson and D. Hill. 2014. "Good Is Not Good Enough: The Benchmark Stroke

Door-to-Needle Time Should Be 30 Minutes." *Canadian Journal of Neurological Sciences* 41 (6): 694–696.
Morgan, L. E., A. C. Titman, D. J. Worthington, and B. L. Nelson. 2016. "Input Uncertainty Quantification for Simulation Models with Piecewise-Constant Non-Stationary Poisson Arrival Processes." In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, 370–381. Piscataway, New Jersey: Institute of Electronic and Electrical Engineers, Inc.
R Core Team. 2016. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. https://www.R-project.org/.
Saver, J. L. 2006. "Time Is Brain - Quantified." *Stroke* 37 (1): 263–266.
Statistics Norway. 2017a. "Table: 07459: Population, by Sex and One-Year Age Groups." https://www.ssb.no/statistikkbanken/SelectVarVal/Define.asp?MainTable=NY3026&PLanguage=1
Statistics Norway. 2017b . "Table: 10826: Population, by Sex and Age." https://www.ssb.no/statistikkbanken/SelectVarVal/Define.asp?MainTable=FolkemBydEttar&PLanguage=1
The AnyLogic Company. 2017. *AnyLogic* (version 8.0.5). www.anylogic.com.
Viana, J., K. Rand-Hendriksen, T. B. Simonsen, M. Barra, and F. Dahl. 2016. "Do Hybrid Simulation Models Always Increase Flexibility to Handle Parametric and Structural Changes?" In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, 1439–1450. Piscataway, New Jersey: Institute of Electronic and Electrical Engineers, Inc.

AUTHOR BIOGRAPHIES

KIM RAND-HENDRIKSEN is a clinical psychologist currently working mostly with statistics and simulation models. He is a senior researcher at the Health Services Research Centre in the Research Department of Akershus University Hospital, and a post-doctoral fellow at the Department of Health Management and Health Economics, University of Oslo. He holds a PhD in health services research, and his interests include Quality Adjusted Life Years, Health Related Quality of Life, Health State Valuation Procedures, the EQ-5D, Behavioral Economics, Cognitive Biases, Validity, Health Psychology, statistics in general, and most other things in particular. His email is kim.rand-hendriksen@ahus.no.

JOE VIANA is a Research Fellow at the Health Services Research Centre in the Research Department of Akershus University Hospital. 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 Joe.Viana@ahus.no.

FREDRIK A. DAHL is a senior researcher at the Health Services Research Centre in the Research Department of Akershus University Hospital. He holds a PhD degree in informatics. He has previously worked with combat simulation for the Norwegian Armed Forces and Bayesian MCMC simulation. His email is Fredrik.A.Dahl@ahus.no.