

LEARNING ABOUT SYSTEMS USING MACHINE LEARNING: TOWARDS MORE DATA-DRIVEN FEEDBACK LOOPS

Mahmoud Elbattah

Owen Molloy

College of Engineering and Informatics
National University of Ireland Galway
University Road, Galway City, IRELAND

ABSTRACT

Machine Learning (ML) has demonstrated great potentials for constructing new knowledge, or improving already established knowledge. Reflecting this trend, the paper lends support to the discussion of why and how should ML support the practice of modeling and simulation? Subsequently, the study goes through a use case in relation to healthcare, which aims to provide a practical perspective for integrating simulation models with data-driven insights learned by ML models. Through a realistic scenario, we utilise ML clustering in order to learn about the system's structure and behaviour under study. The insights gained by the clustering model are then utilised to build a System Dynamics model. Recognizing its current limitations, the study is believed to serve as a kernel towards promoting further integration between simulation modeling and ML.

1 INTRODUCTION

Systems modeling can be contemplated as the science of complexity. As described by (Meadows 2008), systems inherently exhibit adaptive, dynamic, goal-seeking, self-preserving, and sometimes evolutionary behaviour, making a system more than simply the sum of its parts. The process of learning about systems of high complexity can therefore be a challenging task.

From a different perspective, systems can now be dealing with extraordinary amounts of data (i.e. Big Data). It should be taken into account that systems involved with such Big Data scenarios place further burdens on the modeling process, which can go beyond human capabilities in many aspects. For instance, a system's knowledge may need to be learned or extracted from huge amounts of data, which may be accumulating with a high velocity as well. Viewed this way, the complexity of systems can also be interpreted in terms of the complexity of data encompassing the system's knowledge.

In this respect, this study endeavours to investigate the feasibility of utilising Machine Learning (ML) techniques in order to understand, and learn about the underlying structure or behaviour of systems. Our fundamental view is that system models can be developed based on knowledge learned by ML models in tandem with mental models-driven knowledge. The study specifically focused on the System Dynamics (SD) approach, as an exemplary method widely used for systems modeling. Developed by Jay Forrester in the 1950s, the SD approach aimed to introduce a set of tools that enable to understand the structure and dynamics of complex systems. Over decades, SD has successfully established a large community in academia and industry, and continued to become an instrumental artefact for finding appropriate strategies and policies by elucidating the dynamic behaviour of systems.

The paper can be viewed as organised into two main parts as follows. The first part initiates a discussion regarding the process of learning about systems from the perspective of feedback loops. That discussion was aimed to serve as an opening to the rationale behind our arguments. Afterwards, the second part, starting from Section 6, provides a more practical standpoint based on a use case in relation

to healthcare. The use case was mainly adopted in order to present realistic scenarios, where systems modeling can be supported by data-driven insights gained through ML.

2 MOTIVATION: HOW CAN ML ASSIST LEARNING ABOUT SYSTEMS?

In an early insightful analysis, (Sterman 1994) asserted that:

The challenge is how to move from generalisations about accelerating learning and systems thinking to tools and processes that help us understand complexity, and design better policies.

Sterman was notably alluding to the importance of developing assistive tools that can support the process of learning about systems, and understanding their complexity. In this sense, the key impetus for the study was that ML should be further considered as a valid path for that task, especially with the significant momentum gained by data-driven analytics over the past years. To focus the study's objectives, a set of motivational questions were specifically addressed as listed in Table 1. Through the paper we endeavoured to discuss these questions, and provide practical examples that can demonstrate how ML can support the practice of modeling and simulation (M&S).

Table 1: Motivational questions.

Question	Motivation
Q1) How can ML be employed to assist the conceptualisation of a system?	Utilising ML as an assistive tool within the process of learning about the structure or behaviour of systems.
Q2) Is it possible to integrate mental models with ML models in a way that supports the learning process to develop based on a more data-driven manner? If so, how?	The limitations of our mental models raise a need to consider more relatively unbiased reasoning methods for exploring and describing systems, and designing the corresponding models.
Q3) Which ML techniques can be appropriate for the perception of a system's structure, or the behaviour involved within a problem?	Exploring the possible approaches/methods (e.g. supervised or unsupervised learning) to avail of ML for the purpose of learning about systems.
Q4) Can the integration of ML lead to a higher level of confidence in simulation models, indicated by the accuracy of ML models?	The predictive accuracy of ML models can be more measurable. This may in turn extend the confidence in simulation models designed based on insights from ML models in tandem with mental models.

3 RELATED WORK

In an era marked by data-driven knowledge, the M&S community has been reconsidering the emerging opportunities and challenges to the field. For instance, (Taylor et al. 2013) introduced the term "Big Simulation" for describing one of the grand challenges to the M&S research community. Big Simulation describes issues of scale for Big Data input, very large sets of coupled simulation models, and the analysis of Big Data output from these simulations, all running on a highly distributed computing platform. In a thought-provoking study, (Tolk 2015) envisioned that the next generation of simulation models will be integrated with ML, and Deep Learning in particular. The study argued that bringing M&S, Big Data, and Deep Learning together can create a synergy allowing to significantly improving services to other sciences. Similarly, other studies such as (Pruyt 2014) and (Pruyt 2017) interestingly turned the attention to the potentials of integrating SD models with Big Data, or other disciplines related to Data Science.

However, literature obviously lacked pragmatic studies that practically demonstrate the integration of simulation models and ML, to the best of our knowledge. We believe that the M&S community needs much further studies that encourage and popularise that integration, and its potential benefits.

4 BACKGROUND: THE FEEDBACK LOOP CONCEPT

To set a context, we initially intended to drive the discussion through the concept of feedback loop. The feedback concept has been an essential component for the SD approach, and systems modeling in general. This section aims to briefly review how that concept was perceived within the context of systems.

4.1 The Feedback Loop in System Dynamics

The feedback loop concept has its roots in different disciplines that go back further beyond the development of SD. (Richardson 1983) provided a comprehensive historical review of that concept, and how it was endorsed in a number of sciences. The review argued that the feedback concept has evolved and matured as a blend of ideas from different sources of subjects including: i) Engineering, ii) Biology, iii) Mathematical models of biological and social systems, and iv) Social sciences. Richardson interestingly presented examples of feedback loops in engineering that dated back to 250 B.C.

From the perspective of the SD approach, the concept of feedback has been embraced as a fundamental idea within systems modeling. (Forrester 1960; 1961; 1964) kept asserting the importance of feedback loops within systems, and that all decisions are developed within the context of feedback loops. As (Forrester 1968) described it, a feedback loop is a closed path connecting in sequence a decision controlling an action, a state of the system, and information about the state, which returns to the decision-making point. Figure 1(a) sketches the feedback loop in its simplest form.

Further articulations of the feedback loop were compiled by (Sterman 1994). Sterman discussed issues pertaining to feedback loops, and illustrated by examples its significance in a broader context. Equally important, Sterman's work endorsed the inevitable influence of mental models while perceiving information from feedback loops, and making decisions. Figure 1(b) re-portrays the feedback loop in view of existing premises derived from mental models. Furthermore, Figure 1(c) demonstrates that feedback from the real world can also cause changes in our mental models.

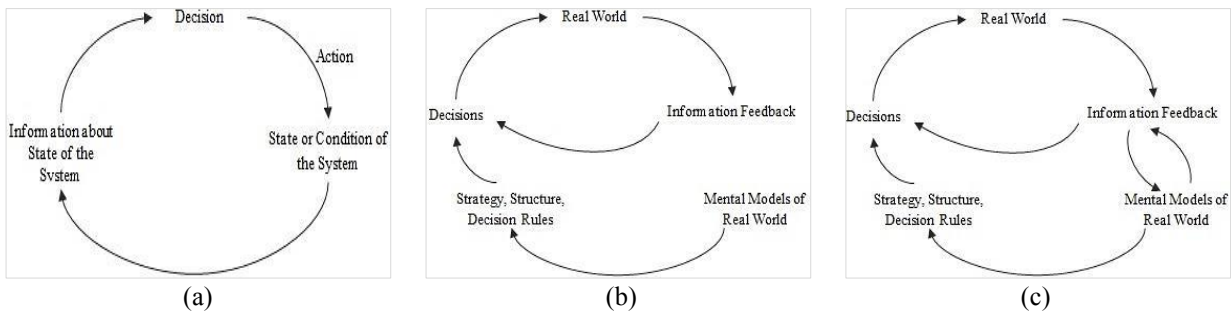


Figure1: The feedback loop concept.

4.2 Limitations of Mental Models

In his landmark textbook (*The Fifth Discipline*), (Senge 1990) described mental models as the deeply ingrained assumptions, generalisations, or even images that affect how we understand the world, and how we take action. From the very beginnings of SD development, (Forrester 1961; 1971) highlighted the unavoidable limitations of our mental models. Forrester emphasised that the mental model is fuzzy, incomplete, and imprecisely stated. However, (Sterman 1994) argued that most people do not appreciate the “ubiquity and invisibility” of mental models. In this regard, Sterman identified a set of mental models-related barriers to learning feedback as follows: i) Misperceptions of feedback, ii) Flawed cognitive maps of causal relations, iii) Erroneous inferences about dynamics, iv) Unscientific reasoning, v) Judgmental errors and biases, and vi) Defensive routines and interpersonal impediments to learning.

It can be understood that the above-mentioned barriers are mostly attributed to the nature of human-based reasoning predicated on biased perception of information. Our initial view was that more machine-oriented assistive methods (e.g. ML) may constitute a key factor to mitigate such limitations.

5 OUR APPROACH: MENTAL MODELS AIDED BY MACHINE LEARNING MODELS

“We are flooding people with information. We need to feed it through a processor. A human must turn information into intelligence or knowledge” – Grace Hopper.

In the sense of that quote by Grace Hopper, one of the early computing pioneers, our approach mainly aimed to support mental models with data-driven knowledge learned by ML. The key idea hinges on the premise that mental models can be assisted by ML models trained to make predictions on a particular aspect of the system’s structure, or behaviour being modelled. It is assumed that changes in the states or conditions of a system can be inferred, at least partially, by ML models.

As illustrated in Figure 2, new data (i.e. feedback) can be generated by new system’s states. Based on data-driven feedback, ML models can be trained to predict the future behaviour of the system. Moreover, ML models can be continuously re-fitted to echo feedback loops, and reflect new system’s conditions. In this manner, the new system state can be learned based on ML models in tandem with mental models. The under-consideration argument is that relatively unbiased, or less biased, data-driven predictions can help improve the understanding of systems, and in turn more accurate decisions can be made.

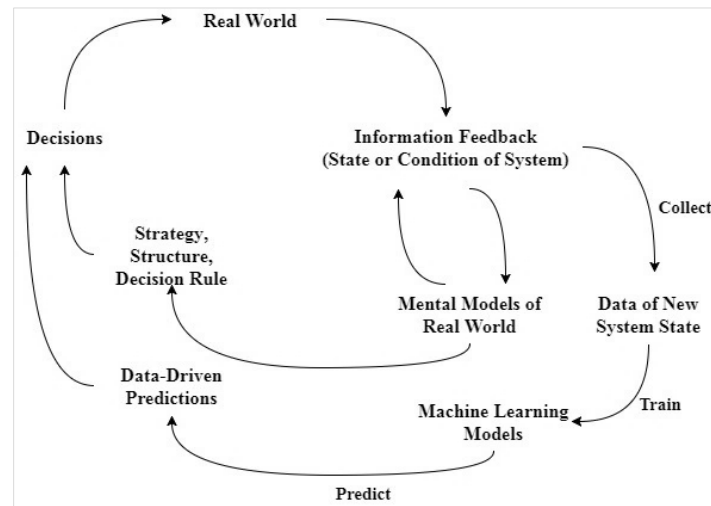


Figure 2: Approach overview.

6 USE CASE: MODELING FLOW OF ELDERLY PATIENTS

For the purpose of demonstrating the applicability of our approach, a case study was developed in relation to healthcare. The following sections elaborate the case setting, and the development of the ML and SD models. The main goal of the use case was to provide a practical scenario where SD models can be designed or adjusted in accordance with new system conditions learned by the aid of ML.

6.1 Case Description

The use case was developed within a healthcare context, with a particular focus on hip fracture care in Ireland. Hip fractures are a major cause of injuries and morbidity among the elderly. As acknowledged by numerous studies (e.g. (Melton 1996)), hip fractures were observed to be exponentially increasing with age, despite the existence of rate variability from country to another. Further, the care of hip fractures has a considerable importance, whereas the Ireland’s Health Service Executive (HSE) identified hip fractures as one of the most serious injuries resulting in lengthy hospital admissions and high costs (HSE 2008).

The typical patient’s journey is described as follows. Initially, a patient is usually received at the ED. The primary surgery is performed after admission to an orthopaedic ward. Subsequently, the patient can possibly undergo various assessments based on falls history and fragility. Eventually, the discharge destination is mainly decided as: i) Home, or ii) Long-stay care facility (e.g. nursing home).

6.2 Data Description

We acquired a dataset from the Irish Hip Fracture Database (IHFD). The IHFD repository is the national clinical audit developed to capture care standards and outcomes for hip-fracture patients in Ireland. The dataset particularly included records about elderly patients aged 60 and over. The data comprised about 8K records over three years (2013-2015). Figure 3(a) plots a histogram of the age distribution of patients, while Figure 3(b) shows the gender percentages. It is noteworthy that a particular patient may be related to more than one record in case of recurrent fractures. However, we were unable to determine the proportion of recurrent cases whereas patients had no unique identifiers, and records were completely anonymised for privacy purposes.

The dataset records contained ample information about the patient's journey. Specifically, a typical patient record included 38 data fields such as gender, age, type of fracture, date of admission and LOS. A thorough explanation of the data fields was available via the official data dictionary (HSE 2015).




Figure 3: The distribution of patients' age and gender in the dataset.

6.3 Purpose of the Model

We aimed to develop an SD model that can depict the flow of elderly hip-fracture patients from admission to discharge. The model can be utilised to understand and estimate the potential demand of care against the capacity of healthcare facilities. Specifically, the model focused on the utilisation of healthcare facilities in terms of: i) Inpatient length of stay (LOS), and ii) Discharge destinations. To focus the purpose of the model, the questions of interest are stated as below:

1. What is the expected consumption of hospital resources with regard to the inpatient LOS?
2. What is the expected proportion of elderly patients discharged to home, or long-stay care?

6.4 Selection of Machine Learning Technique

As per the questions posed in Section 2, we initially aimed to investigate the possible ML techniques that can be used to assist the process of problem conceptualisation. In this regard, unsupervised ML techniques (e.g. clustering) present adequately for perceiving the system's structure or behaviour as follows. The SD models typically deal with aggregate entities (e.g. population of patients), and not individual entities, or agents. Those aggregate entities are represented as "stocks", which characterise the state of the system and generate the information upon which decisions and actions are made (Sterman 2000). In a relatively similar manner, clustering seeks to realise the segmentation of a heterogeneous population into a number of more homogeneous subgroups (Aldenderfer and Blashfield 1984). Viewed this way, the suggested homogeneous groups (i.e. clusters) may correspond to particular stocks in the SD model. In addition, clustering is an effective method for exploring potential underlying structures in the system without making any prior assumptions that might be biased.

6.5 Discovering Patient Clusters

As explained, ML clustering techniques can hold potentials for learning about systems. In this respect, the study employed clustering to realise the segmentation of patients from a data-driven viewpoint. The following sections explain the data pre-processing procedures, and clustering experiments.

6.5.1 Data Pre-processing

i) Outliers Removal: According to (NOCA 2014), the mean and median LOS for hip- fracture patients were reported as 19 and 12.5 days respectively. Therefore, we considered LOS values longer than 60 days as outliers, which represented about 5%. Figure 4 plots the LOS histogram.

ii) Feature Scaling: Several studies (e.g. (Patel and Mehta 2011)) argued that large variations within the range of feature values can affect the quality of computed clusters. Therefore, the features were rescaled to a standard range using the min-max normalisation method.

iii) Feature Extraction: In a report published by (British Orthopaedic Association 2007), six quality standards for hip fracture care were emphasised. Those standards generally reflect good practice at key stages of care including: i) All patients should be admitted to an orthopaedic ward within 4 hours, and ii) All patients should have surgery within 48 hours of admission. The raw data did not include fields that explicitly captured such standards. However, they could be derived based on the date-time values of patient arrival, admission and surgery. In this way, two new features were added named as “Time to Admission (TTA)” and “Time to Surgery (TTS)”. However, only the TTS was eventually included, whereas the TTA contained a significant amount of missing values.

iv) Feature Selection: The K-Means algorithm is originally applicable to numeric features only, where a distance metric (e.g. Euclidean distance) can be used for measuring similarity between data points. For this reason, we considered the numeric features only. Specifically, the model was trained using the following features: i) LOS, ii) Age, and iii) TTS.




Figure 4: The distribution of inpatient LOS within the dataset.

6.5.2 Clustering Experiments

The partitional clustering approach was embraced using the widely prevalent K-Means algorithm. The problematic question while approaching a clustering task is how many clusters (K) may exist? In our case, the number of clusters was experimented with K ranging from 2 to 7. Initially, the quality of clusters was inspected based on the within cluster sum of squares (WSS) (Figure 5). In light of that, it could be initially suggested that three or four clusters of patients can likely exist within the dataset.

In order to determine the most appropriate number of clusters, the suggested clusters were projected into two dimensions based on the Principal Component Analysis (PCA), as shown in Figure 6. Each sub-figure below represents the output of a clustering experiment using a different K. Initially with K=2, the output indicated a promising tendency of clusters, where the data space is obviously separated in two big clusters. Similarly for K=3, the clusters were still well-separated. However, the quality of clusters started to decline when K=4 onwards. Eventually, it turned out that there were three clusters that best separated the dataset into coherent patient cohorts. We availed of the Azure ML Studio to train the clustering model. The cluster visualisations were produced by the R-package ggplot2 (Wickham 2009).




Figure 5: The sum of squared distances within clusters.

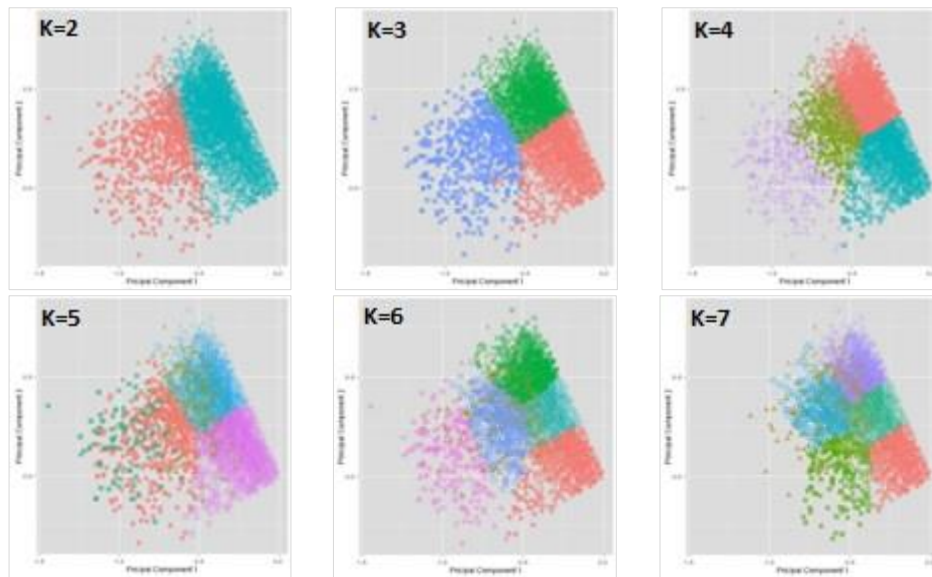


Figure 6: Visualisation of clustering experiments with the number of clusters (K) ranging from 2 to 7.

6.5.3 Learning Data-Driven Insights from Clusters

In this section, we start exploring the discovered clusters mostly through visualisations. Our intention was to reveal potential correlations or insights, which can assist with the SD model design. The clusters were examined with respect to patient characteristics (e.g. Age), outcomes (e.g. LOS), and important care-related factors (e.g. TTS). In Figure 7(a), the inpatient LOS is plotted with respect to the three discovered clusters of patients. At first glance, it is obviously observable that the patients of Cluster3 had a much longer LOS rather than Cluster1 and Clusters2. On the other hand, Cluster1 and Cluster2 shared a very similar distribution of the LOS variable, apart from a few outliers in Cluster2.

Second, we examined the clusters with respect to the TTS. As before-mentioned, the TTS is considered as one of the quality standards for hip fracture care. Once again, the patients of Cluster3 were observed for having a relatively longer TTS than those patients of Cluster1 and Cluster2. Likewise the LOS, Cluster1 and Cluster2 experienced a quite similar distribution of the TTS. Figure 7(b) plots the TTS variation against the three clusters of patients.

For its considerable emphasis within elderly care, the clusters were also explored regarding the patient age. In our context, the possibility of sustaining hip fractures can increase significantly by ageing. It turned out that Cluster1 and Cluster3 tended to have relatively older patients rather than Cluster2. Figure 7(c) plots the age distribution within the three clusters. Eventually, the clusters were inspected with regard to discharge destinations. Discharge destinations can be generally classified into: i) Home, or i) Long-stay care facility. Patients discharged to long-stay care can likely spend prolonged periods of residential care. As per Figure 8, there was a pronounced variation between the clusters in this regard.

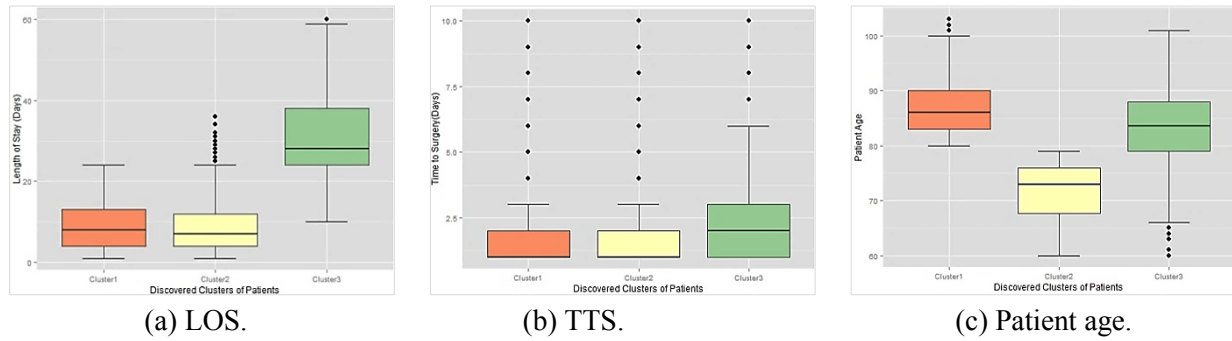


Figure 7: The variation of the LOS, TTS, and age variables in the patient clusters.




Figure 8: The variation of discharge destinations in the patient clusters.

6.6 SYSTEM DYNAMICS MODELING

6.6.1 Initial SD Model

The initial model provided a bird's-eye view of the care scheme of hip fracture with respect to the questions of interest. The model focused on capturing the dynamic behaviour in relation to the continuous growth of ageing, and the consequent implications on the incidence of hip fractures among the elderly. The main actors within the model were defined as follows: i) Elderly patients, ii) Acute hospital, and iii) Discharge destinations including home or long-stay care facilities.

Two different inflow rates were used for male and female patients, as defined by (Dodds, Codd, Looney and Mulhall 2009). The model included a single reinforcing loop implied by the elderly patients of a fragility history, who are susceptible to re-sustain hip fractures, or fall-related injuries at least. At this stage, the model did not consider the different characteristics of patients learned by the ML clustering experiments. Figure 9 illustrates the initial SD model.

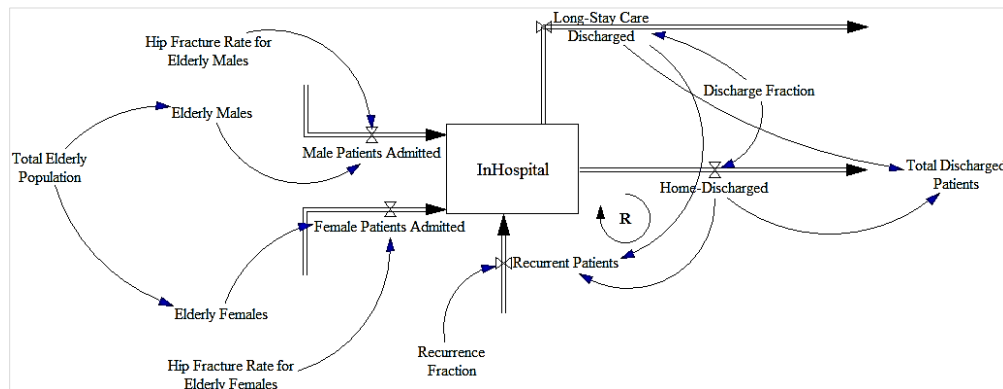


Figure 9: Initial SD model.

6.6.2 Cluster-Based SD Model

The SD model was re-designed in light of the knowledge learned by clustering experiments. In particular, the model was disaggregated into 3 different stocks representing the discovered clusters of patients. Furthermore, the model behaviour was mainly set based on the cluster analysis. For instance, the first and second clusters were considered to undergo the same TTS delay (i.e. TimeToSurgery1), while the third cluster was set a different delay (i.e. TimeToSurgery2). Likewise, each cluster was associated with a specific discharge destination-related fraction.

Equally important, the inflow of elderly patients was structured based on the age groups within the clusters. In particular, both of the first and third clusters were modelled to contain more elderly patients (i.e. aged 80-100), while the second cluster was associated with less elderly patients (i.e. aged 60-80). This reflected the age distribution within the clusters, as sketched previously in Figure 7(c). For the purpose of simplicity, the model did not include the case of recurrent patients, which caused a reinforcing loop in the initial model. Figure 10 sketches the cluster-based model.

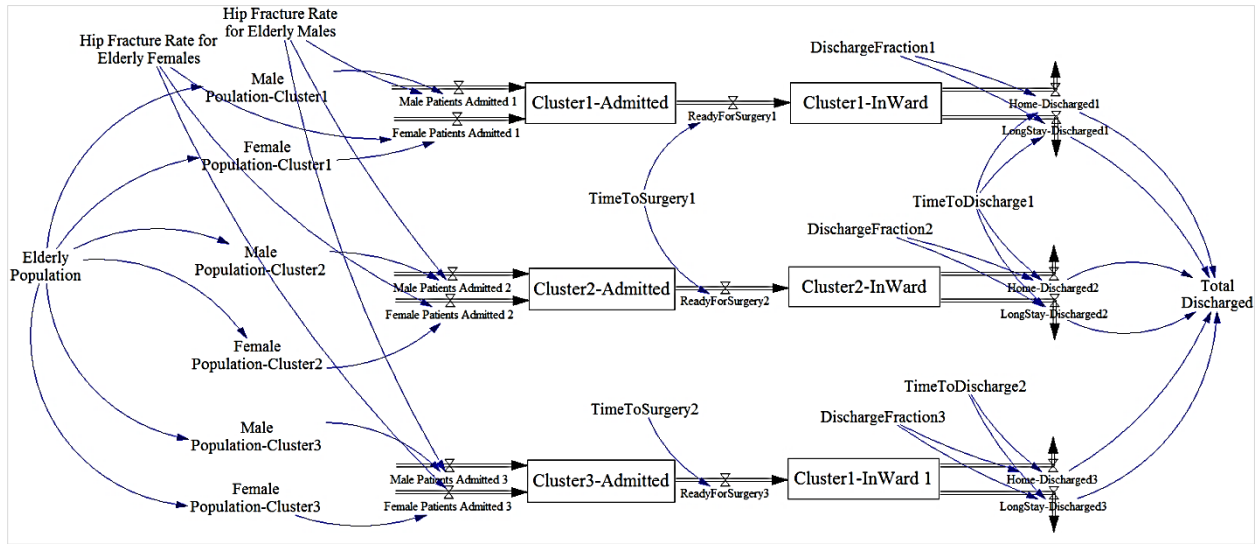


Figure 10: The cluster-based SD model.

6.6.3 Simulating Data-Driven Feedback

In order to demonstrate the effect of data-driven feedback learned by ML, we applied a hypothetical scenario of care improvement. The scenario was intended to simulate a hypothetical change in the system behaviour as follows. It was assumed that a new policy was introduced starting from the year 2014 towards improving the patient's journey. The new policy aimed to maintain the hip-fracture care standards by keeping the TTA and TTS within 4 hours and 48 hours respectively. In accordance with the new policy, the average inpatient LOS was assumed to decrease by 20% and 30% in 2014 and 2015 respectively. Further, the proportion of patients discharged to long-stay residential care was assumed to decrease by 5% and 10% in 2014 and 2015 respectively. In order to reflect the new policy, the patient records of the years 2014 and 2015 were synthetically altered. For instance, the LOS was reduced by 20% for patients discharged in 2014.

Subsequently, the clustering model was retrained in view of the policy changes. The new clusters are demonstrated in Figure 11. It turned out that the new policy led to fewer clusters of patients. Specifically, the finest separation of clusters was realised when $K=2$. The new clusters were re-explored with respect to the LOS, TTS, and patient age as plotted in Figure 12. Based on the new patient clusters, the SD model was re-designed. The updated SD model corresponded to knowledge updates learned by the ML clustering model. Figure 13 sketches the updated SD model.

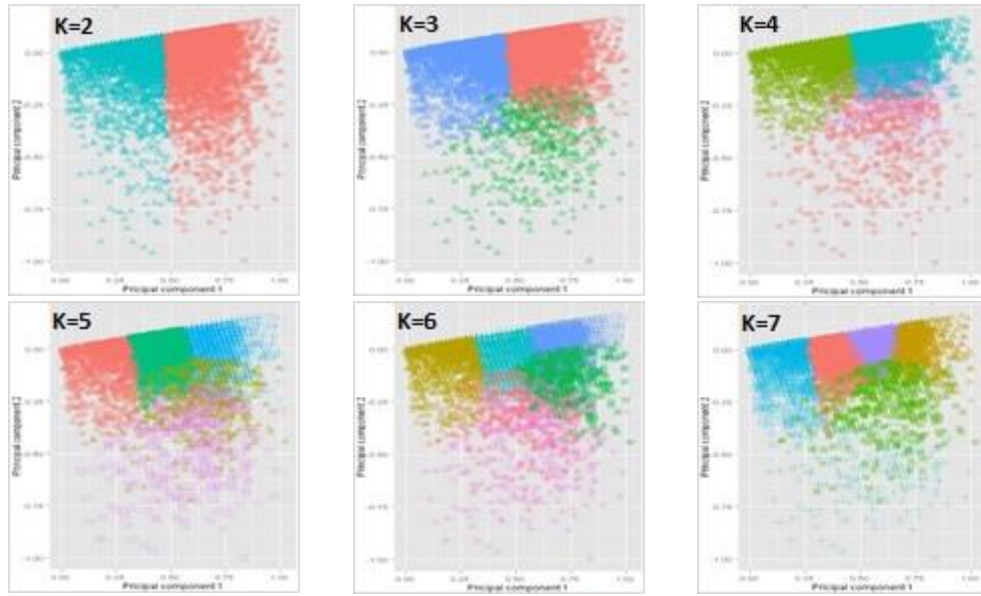


Figure 11: Visualisation of clustering experiments after applying the new care policy.

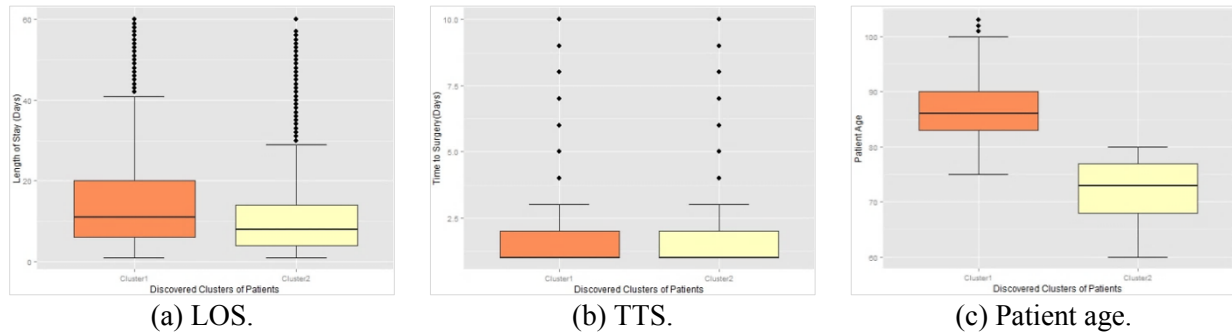


Figure 12: The variation of the LOS, TTS, and age variables in the new clusters.

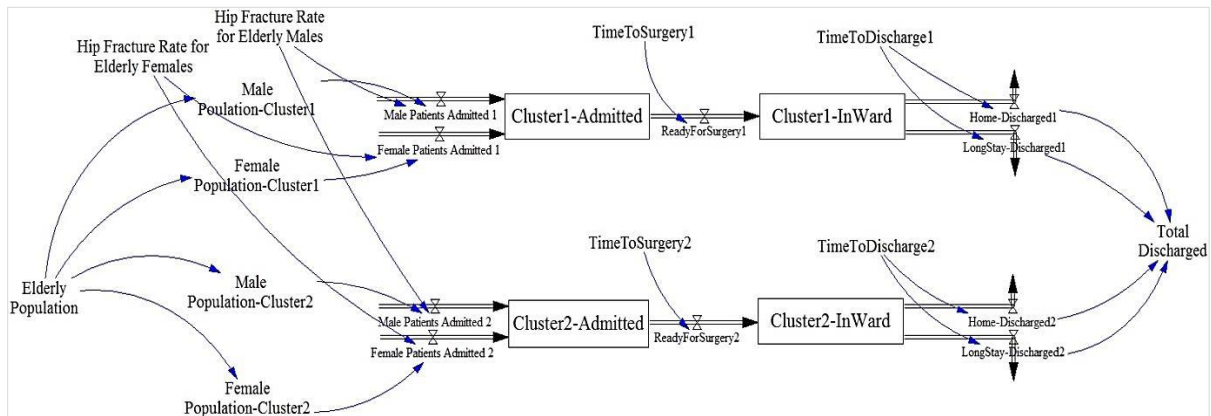


Figure 13: The updated cluster-based model.

7 DISCUSSION

We believe that the developed scenario largely addressed the motivational questions listed in Table1. First, the clustering model was employed effectively for the purpose of understanding the system structure, where the SD model stocks actually represented the three discovered clusters of patients.

Moreover, the variations within the clusters in terms of patient characteristics (e.g. age), or care-related factors (e.g. TTS) assisted with shaping the model behaviour. Furthermore, it can be argued that the SD model was constructed with an established confidence based on the clustering model. The well-validated quality of clusters along with the compelling visualisations could support the rationale behind the SD model design in terms of structure and behaviour as well. Thus, the use of ML could have led to lowering the epistemic uncertainty usually attributed to the subjective interpretation of system knowledge by modelers, or simulationists, as explained by (Oberkamp 2002).

In our case, the clustering model played an appropriate role while trying to explore possible systemic structures based on a pure data-driven standpoint. However, other ML techniques may be more appropriate within different situations, or other simulation approaches.

8 STUDY LIMITATIONS

A set of limitations are acknowledged as follows. The presented use case may not have been the best exemplary scenario to demonstrate the potentials of integrating simulation modeling and ML. We believe that a typical Big Data scenario can better present the benefits of that integration.

Another relevant issue of concern, the patient clustering was based on a mere data-driven standpoint. Adding a clinical perspective (e.g. diagnosis, procedures) may group patients differently.

9 CONCLUSIONS AND FUTURE DIRECTIONS

The integration of mental models with data-driven insights learned by Machine Learning (ML) models can yield potential benefits for the practice of modeling and simulation. One benefit is lowering the bias of mental models, which can in turn increase the confidence in simulation models. In this regard, the study attempted to practically demonstrate how ML can assist with building simulation models.

Looking into the future, more sophisticated ML techniques can be effectively used in order to distil the knowledge underlying further complex systems. For example, it would be interesting to investigate how simulation models can be integrated with Deep Learning (DL). DL (LeCun, Bengio and Hinton 2015) received wide attention within the ML research for its capacity that dramatically improved the state-of-the-art in hard problems such as visual object recognition, or speech recognition. Using multiple processing layers, DL allowed models to learn data representations with multiple levels of abstraction.

In the systems world, there are similar complex problems that might be quite intractable to be analytically described or perceived (e.g. biology systems). We conceive that Deep Neural Networks (DNN) can be utilised to help understand such complex systems, and better predict their behaviour. Furthermore, the DNN can be incrementally trained on a timely basis by the arrival of new data that echo new states of the system. In this manner, the DNN training can capture and update the system's knowledge in a semi-automated manner. This can extend further opportunities for modeling dynamic systems that inherently exist within rapidly changing environments.

REFERENCES

- Pruyt, E., 2017. "Integrating Systems Modelling and Data Science: The Joint Future of Simulation and 'Big Data' Science". In *Artificial Intelligence: Concepts, Methodologies, Tools, and Applications* (822-840). IGI Global.
- Forrester, J.W. 1961. *Industrial Dynamics*. MIT Press, Cambridge, MA.
- Forrester, J.W. 1960. "The Impact of Feedback Control Concepts on the Management Sciences". Foundation for Instrumentation Education and Research.
- Forrester, J.W. 1964. "Common Foundations Underlying Engineering and Management". *IEEE spectrum*, 1(9), pp.66-77.
- Richardson, G. 1983. "The Feedback Concept in American Social Science, with Implications for System Dynamics." In *Proceedings of the 1983 International System Dynamics Conference*, Massachusetts.

- Sterman, J. D. 1994. "Learning in and about Complex Systems". *System Dynamics Review*, 10(2-3).
- Forrester, J.W. 1968. *Principles of Systems*. MIT Press, Cambridge, MA.
- Senge, P. 1990. *The Fifth Discipline: The Art and Science of the Learning Organization*. New York: Currency Doubleday.
- Oberkampff, W.L., DeLand, S.M., Rutherford, B.M., Diegert, K.V. and Alvin, K.F., 2002. "Error and Uncertainty in Modeling and Simulation". *Reliability Engineering & System Safety*, 75(3).
- Aldenderfer, M.S. and Blashfield, R.K., 1984. *Cluster Analysis*. Sage Publications, Beverly Hills, CA.
- NOCA. 2014. "Irish Hip Fracture Database National Report 2014". <https://goo.gl/R18SX2>
- Patel, V.R. and Mehta, R.G., 2011. "Impact of Outlier Removal and Normalization Approach in Modified K-Means Clustering Algorithm". *International Journal of Computer Science Issues*, 8(5).
- Sterman, J.D. 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. McGraw-Hill Publishing, Boston, USA.
- Forrester, J.W. 1971. "Counterintuitive Behavior of Social Systems". *Technological Forecasting and Social Change*, 3, 1-22.
- Dodds, M.K., Codd, M.B., Looney, A. and Mulhall, K.J., 2009. "Incidence of Hip Fracture in the Republic of Ireland and Future Projections: A Population-Based Study". *Osteoporosis International*, 20(12), pp.2105-2110.
- Pruyt, E., Cunningham, S., Kwakkel, J.H. and De Bruijn, J.A., 2014. "From Data-Poor to Data-Rich: System Dynamics in the Era of Big Data". In *Proceedings of the 2014 International Conference of the System Dynamics Society*, System Dynamics Society, Delft, Netherlands.
- Tolk, A., 2015, July. "The Next Generation of Modeling & Simulation: Integrating Big Data and Deep Learning". In *Proceedings of the Conference on Summer Computer Simulation* (pp. 1-8). International Society for Computer Simulation.
- Melton, L. J. 1996. "Epidemiology of Hip Fractures: Implications of the Exponential Increase with Age." *Bone* 18(3), 121-125.
- British Orthopaedic Association, 2007. *The Care of Patients with Fragility Fracture*. London: British Orthopaedic Association, pp.8-11.
- HSE. 2008. "Strategy to Prevent Falls and Fractures in Ireland's Ageing Population".
- Meadows, D.H. 2008. *Thinking in Systems: A primer*. Chelsea Green Publishing, Vermont, USA.
- Wickham, H., 2009. *ggplot2: Elegant Graphics for Data Analysis*. Springer, New York, USA.
- HSE. 2015. "HIPE-Data Dictionary". <https://goo.gl/YTeqjH>.
- LeCun, Y., Bengio, Y. and Hinton, G., 2015. "Deep Learning". *Nature*, 521(7553), pp.436-444.
- Taylor, S.J., Khan, A., Morse, K.L., Tolk, A., Yilmaz, L. and Zander, J., 2013, April. "Grand Challenges on the Theory of Modeling and Simulation". In *Proceedings of the Symposium on Theory of Modeling & Simulation-DEVS Integrative M&S Symposium* (p. 34). Society for Computer Simulation International.

AUTHOR BIOGRAPHIES

MAHMOUD ELBATTAH is a PhD candidate at the National University of Ireland, Galway. He is a holder of the Hardiman Scholarship from the College of Engineering and Informatics. His research interests are in the fields of simulation modeling, and machine learning. His e-mail address is m.elbattah1@nuigalway.ie.

OWEN MOLLOY is a Lecturer in Information Technology at the National University of Ireland, Galway. His research interests are in the areas of business process modelling and simulation, software engineering, healthcare processes and information systems. His e-mail address is owen.molloy@nuigalway.ie.