

AI to enhance interactive simulation-based training in resuscitation medicine

Rob Brisk¹, Raymond Bond², Jun Liu², Dewar Finlay³, James McLaughlin³ and David McEneaney¹

¹ Cardiovascular Research Unit, Craigavon Area Hospital, 68 Lurgan Rd, Portadown, BT63 5QQ

² Computer Science Research Institute, Ulster University, Shore Road, Newtonabbey, BT37 0QB

³ Nanotechnology and Integrated BioEngineering Centre, Ulster University, Shore Road, Newtonabbey, BT37 0QB

Correspondence to robbrisk@hotmail.com

When patients become acutely unwell, the ability of frontline healthcare professionals to act quickly and effectively can mean the difference between life and death. High-fidelity simulation is the gold standard by which medics acquire and maintain key resuscitation skills, but the resource-intensive nature of current, face-to-face training limits access to training and allows “skills fade” to creep in. We propose that human computer interaction-based simulations augmented by artificial intelligence could provide a cost-effective alternative to traditional training and allow clinicians much greater access to training. This paper is mostly an in-depth discussion; however, we also present a 3D simulator for resuscitation skills training which we developed using the Unity games physics engine.

Digital simulation. Resuscitation medicine. Artificial intelligence. Reinforcement learning.

1. INTRODUCTION

Much is made of the potential of emerging artificial intelligence (AI) technology to bring badly-needed innovation to the field of medicine, and yet currently we see little evidence of this on the wards each day. Part of the problem might be that the vanguard of this march of progress is comprised almost exclusively of data scientists and machine learning (ML) researchers, while most healthcare workers remain entirely ignorant of even the basic concepts underpinning ML and, by extension, the technology we commonly define as being “artificially intelligent”. Naturally, practitioners of ML are likely to gravitate towards clinical problems that present favourable targets for their science, for example single-step classification tasks in data-rich areas; hence, disciplines like radiology are enjoying the lion’s share of the attention from the ML community [1].

Here, we propose to explore the application of ML to a sequential decision-making task in a high impact but relatively data-poor area of healthcare, albeit – for now – within the context of a training application rather than a system interacting directly with patients.

2. THE CLINICAL NEED

There are over 10, 000 in-hospital cardiac arrests annually in the UK [2]. Outcomes for these patients are poor: only one in five will survive to hospital discharge and over half of these survivors will have some degree of neurological (brain) damage [3]. Some cardiac arrests happen “out of the blue”, due to sudden events such as myocardial infarction or

pulmonary embolus, but a significant proportion will be preceded by a gradual deterioration in the patient’s condition. It has been concluded that as many as 5% of hospital deaths may be averted, largely by the prompt identification and effective treatment of acute illness [4,5,6,7]. The key question then is: how do we improve the recognition and treatment of the deteriorating patient?

There are a range of novel technology-based solutions on the market but their efficacy remains unproven [8]; high-quality simulation training for clinical staff is still by far the best-evidenced intervention [9,10]. However, the resource-intensive nature of existing, face-to-face simulation methods is a limiting factor in arranging high frequency resuscitation training, due largely to the requirement for a high ratio of expert instructors to trainees. There is evidence to suggest that the optimal training frequency might be as often as six-weekly [11], but, in ever-shorter-staffed healthcare systems, even the logistic challenge of ensuring practitioners have access to an Advanced Life Support (ALS) course just once every four years has necessitated a push by the European Resuscitation Council to streamline training and cut courses from two days to one [12,13]. It is to this



Figure 1: High-fidelity, face-to-face clinical simulation
Reproduced with permission. © University of Dundee

issue that we propose to explore a novel human computer interaction-based solution, in the form of an AI-supported digital simulation system that – by negating the need for an expert human presence – could facilitate the delivery of low-cost, high-impact training at unprecedented frequency.

3. DIGITAL RESUSCITATION SIMULATION

The rationale for digital simulation in clinical training is well established. In fact, for certain procedural skills, such as those required to perform laparoscopic surgery, it has proven more efficacious even than conventional training methods [14]. To explain why, then, we are proposing the need for an AI-based approach and billing this as novel research, it is necessary to first point out that existing simulators cater mostly for procedural skills rather than clinical decision-making skills, and that the two applications require very different educational frameworks.

At the simplest level, gated progression through a surgical simulation, where one particular method or technique is usually considered to be optimal at each step of the procedure, can be achieved using single-condition “if-then-else” statements: if [*the trainee performs step A according to the optimal method*] then [*the trainee is deemed to have demonstrated proficiency and can progress to step B*] else [*they receive constructive feedback and retry step A*]. This is both educationally viable (it allows for the integration of a proficiency-based progression model [15]) and computationally favourable: by thus restricting the permissible action space, the number of resultant states for which the simulation must account is very limited. Naturally, modern simulators have built upon this basic framework to develop less obviously linear narratives and to account for a number of common procedural complications, but by continuing to restrict permissible user actions they can continue to limit the state space to manageable dimensions.

Our proposed resuscitation simulator cannot take advantage of the same approach. The focus when training for ALS moves from procedural to conceptual knowledge [16], because ALS providers are not required to become expert at tackling a fixed problem like their surgical colleagues but rather to develop cognitive processes that are generalizable to a wide range of disparate clinical scenarios, within which they must make prompt and effective decisions. Anyone who is familiar with model-free reinforcement learning will understand that development of generalizable behaviour policies first requires exploration of the action-state space [17], and the same is true for humans – though psychologists would more likely term this as the “active experimentation” phase of Kolb’s experiential learning cycle [18] and in lay terms it is simply known as trial-and-error learning. Thus, a simulator designed to develop generalizable

behaviour is likely to maximise its efficacy by allowing users access to action-state spaces that reflect the true diversity of real world experience.

With this in mind, it is our conclusion that stochastic simulation will offer the best training environment for resuscitation medicine, as it most accurately recreates the unpredictable nature of human physiology (or, at least, the unpredictability resulting from our incomplete understanding thereof) and allows for exposure to a much greater number of possible states than a deterministic framework; indeed, we adopted a stochastic approach for our prototype simulator [19]. However, the sheer volume of dynamic, interdependent variables involved in a high-fidelity resuscitation simulation makes it impractical to hard-code a generalizable ruleset by which to evaluate trainees’ actions for any given state (our prototype simulator used fuzzy logic for this, but as we increased complexity by adding new clinical features and scenarios, such an approach became nonviable), and having a large state space unfortunately compounds the problem by rendering the use of expert consensus regarding the optimal action for individual states – the approach favoured by many surgical simulations – similarly unworkable.

So this leaves us with a problem that has historically limited the value of digital simulation training within many areas of medical practice: how does one devise a system for automatically evaluating clinical performance within a complex, stochastic simulation in order to provide real-time, corrective feedback (without which the simulator is of questionable educational value [20])?

4. THE ROLE OF MACHINE LEARNING

At first glance, supervised learning appears to represent a promising solution here: the data that informs clinical decision-making is readily expressed as a feature vector, and mapping this to the corresponding clinical actions can be reframed as a multinomial classification problem. So where we are struggling to capture the complexity of the relevant medical knowledge base in a generalizable ruleset using conventional computing strategies, a suitable classification algorithm could



Figure 2: Digital simulation for laparoscopic surgery
Reproduced with permission. © Marcus Rall

discern such a ruleset for us.

The idea is an appealing one, but of course is rendered hopelessly unrealistic by the almost complete absence of comprehensive, large-scale datasets in this and many other clinical fields. Furthermore, in resuscitation medicine it is not merely the case that we have been failing to capture data in an appropriate, centralised digital format and that we can simply redress this moving forwards; recording and using data from patients who are often too unwell to give consent is a process fraught with ethical problems and confounded by the fact that, in the heat of a medical emergency, effective data capture is the last thing on anyone's mind. Furthermore, the very rationale for this project is that resuscitation medicine is often performed sub-optimally, thus one might question whether imitation learning is necessarily the best approach at all.

There is, however, an alternative and potentially abundant source of clinical data within our particular problem: the simulator itself.

Although the data derived from the simulator is synthesised rather than captured from real-world practice, it is closely informed by clinical expertise and there is a growing body of evidence to support the efficacy of data synthesis in addressing ML tasks for data-poor areas (particularly within the field of computer vision) [21,22,23]. Furthermore, we are not currently proposing to deploy our ML model in real-world practice, so one could argue that simulator-derived data is the most appropriate training material for our present purposes.

The approach we propose takes its precedent largely from the seminal work by researchers at DeepMind in the field of deep reinforcement learning [24,25,26]. In the initial 2013 study, they employed a deep Q-learning strategy to attain human level performance in three of six complex reinforcement learning tasks involving Atari games. Their system, in short, consisted of a deep neural network tasked with predicting the action-value ("Q") function for a given behaviour policy (usually referred to as " π "). The network was updated during the training process using stochastic gradient descent, and the update process smoothed out using an experience replay mechanism (to avoid, say, a promising behaviour

strategy being too heavily penalised for a single bad outcome). In 2015, further refinements to this process allowed them to surpass human performance in a large number of similar tasks.

This approach works well for environments like the game Space Invaders, where the action-state space is limited and there is minimal need for long-term planning. The ML model can develop generalizable skills quickly and rapidly transition away from an "epsilon greedy" strategy, whereby it spends more time exploiting its new skillset and less time exploring its environment (or, as we described it earlier, engaging in trial-and-error learning).

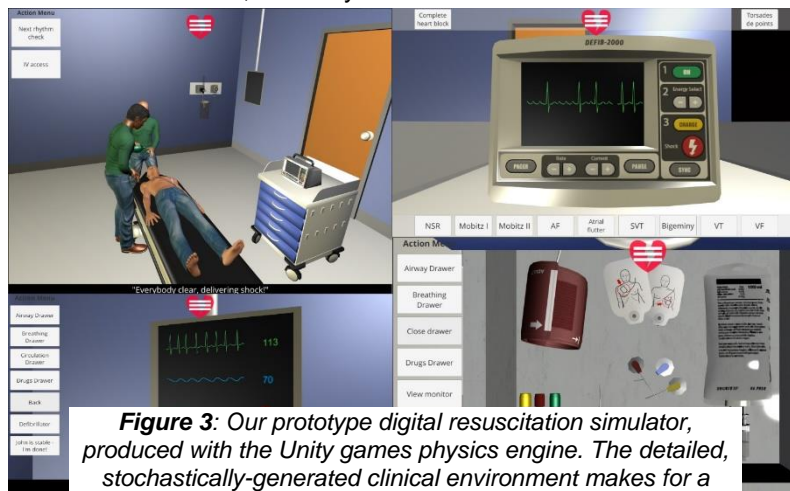


Figure 3: Our prototype digital resuscitation simulator, produced with the Unity games physics engine. The detailed, stochastically-generated clinical environment makes for a particularly high-fidelity experience but necessitates a novel approach to automated trainee evaluation.

Our simulator, however, has a comparatively high-dimensional action space, a more diverse state space, and a greater need for long-term planning, so a ML model functioning therein would need both a much longer period of exploration (or "a slower epsilon decay") and less frequent policy

updates to achieve a similar level of efficacy

using Q-learning, which would result in exponentially increased computational cost. Furthermore, the stochastic nature of the simulator may confound attempts to learn an effective action-value policy (because the same action taken in the same state could potentially result in two different "rewards").

Our aim, therefore, is to take a further lead from the DeepMind researchers: in 2016, they revisited some of the Atari problems, but this time using an "actor-critic" approach. Instead of trying to learn an optimal action-value function (from which the behaviour policy is then implied in a straight Q-learning approach), the actor-critic method employs two asynchronous models: an "actor" whose task is to directly learn an optimal behaviour policy and a "critic" whose task is to learn an action-value policy upon which the actor bases its updates. This has a few key advantages over the Q-learning framework – namely, direct policy optimisation allows the model to more effectively deal with high-dimensional action spaces and to learn stochastic policies, and the addition of a critic model as an action-value estimator offsets the increased variance and allows for more frequent policy updates than, say, a Monte Carlo approach to

policy optimisation [27]. For this reason, we believe the actor-critic framework represents the most promising solution to our particular problem.

5. CONCLUSION AND FUTURE WORK

Through the methods reviewed above, we propose to train an actor-critic model for the task of resuscitating virtual patients within our simulator. The intention would then be to employ the “critic” network from such a model as a means of evaluating the actions of our human trainees within any given state of the simulation, thus providing a basis upon which to provide constructive, real-time feedback within a complex, stochastic clinical simulation.

If our proposed approach is successful, it may solve the problem of delivering high-frequency resuscitation simulation training within a resource-constrained healthcare system, and plausibly improve patient outcomes as a result. However, it could also open up a new paradigm of digital medical education: imagine if final year medical students could hone their clinical skills on simulated wards – receiving constructive feedback as they assessed and treated virtual patients – before they ever made a management decision regarding a real-world patient.

Furthermore, a successful outcome from our research would lay the ground for an exploration of whether this framework – i.e. clinical simulation as a training environment for reinforcement learning models – represents a potential means of creating clinical AI systems that can undertake sequential decision-making tasks directly affecting patient care even in data-poor areas of practice (perhaps using transfer learning strategies to further train the “actor” element of our actor-critic model with whatever real-world data is available).

6. REFERENCES

- 1 Wang S, Summers RM. Machine learning and radiology. *Medical image analysis*. 2012; 16(5):933-51
- 2 Nolan JP, Soar J, Smith GB, et al. Incidence and outcome of in-hospital cardiac arrest in the United Kingdom National Cardiac Arrest Audit. *Resuscitation*. 2014;85(8):987-92.
- 3 Young GB. Clinical practice. Neurologic prognosis after cardiac arrest. *N Engl J Med*. 2009;361(6):605-11.
- 4 National Confidential Enquiry into Patient Outcome and Death. An acute problem? London: NCEPOD; 2005.
- 5 National Institute for Health and Clinical Excellence. Acutely Ill Patients in Hospital: Recognition and Response to Acute Illness in Hospital. National Institute for Health and Clinical Excellence, London. 2007.
- 6 Thomson R.T., Leuttel D., Healey F. & Scobie S. Safer Care for the Acutely Ill Patient: Learning from Serious Incidents. National Patient Safety Agency, London. 2007.
- 7 Hogan H, Healey F, Neale G, Thomson R, Vincent C, Black N. Preventable deaths due to problems in care in English acute hospitals: a retrospective case record review study. *BMJ quality & safety*. 2012; 21(9):737-45.
- 8 Amorim FF, Santana AN. Automated early warning system for septic shock: the new way to achieve intensive care unit quality improvement? *Annals of translational medicine*. 2017; 5(1):17.
- 9 Callaway CW, Soar J, Aibiki M, et al. Part 4: Advanced Life Support: 2015 International Consensus on Cardiopulmonary Resuscitation and Emergency Cardiovascular Care Science With Treatment Recommendations. *Circulation*. 2015;132(16 Suppl 1):S84-145.
- 10 Bellomo R, Goldsmith D, Uchino S, et al. Prospective controlled trial of effect of medical emergency team on postoperative morbidity and mortality rates. *Crit Care Med*. 2004;32(4):916-21.
- 11 Sutton RM, Niles D, Meaney PA, et al. Low-dose, high-frequency CPR training improves skill retention of in-hospital pediatric providers. *Pediatrics*. 2011; 128(1):e145-51.
- 12 G.D. Perkins, P.K. Kimani, I. Bullock, et al. Improving the efficiency of advanced life support training: a randomized. *Controlled Trial Ann Intern Med*. 2012;157:19-28
- 13 G.D. Perkins, P.K. Kimani, I. Bullock, et al. Improving the efficiency of advanced life support training: a randomized. *Controlled Trial Ann Intern Med*. 2012;157:19-28
- 14 Seymour NE, Gallagher AG, Roman SA, et al. Virtual reality training improves operating room performance: results of a randomized, double-blinded study. *Ann Surg*. 2002; 236(4):458-63.
- 15 Gallagher AG, Ritter EM, Champion H, et al. Virtual reality simulation for the operating room: proficiency-based training as a paradigm shift in surgical skills training. *Ann Surg*. 2005;241(2):364-72.
- 16 Schneider M, Rittle-Johnson B, Star JR. Relations among conceptual knowledge, procedural knowledge, and procedural flexibility in two samples differing in prior knowledge. *Developmental psychology*. 2011; 47(6):1525-38.
- 17 R. Sutton and A. Barto. Reinforcement Learning: An Introduction. MIT Press, Cambridge, MA, 1998.
- 18 Kolb, D. A. 1984. *Experiential learning: Experience as the source of learning and development*. New Jersey: Prentice-Hall.
- 19 Brisk, R. *Virtu-ALS*. [COMPUTER / MOBILE APPLICATION]. Available at www.virtu-ALS.com. [Accessed 15 March 2018]
- 20 Hewson MG, Little ML. Giving feedback in medical education: verification of recommended techniques. *J Gen Intern Med*. 1998;13(2):111-6.
- 21 Pishchulin, L., Jain, A., Wojek, C., Andriluka, M., Thormählen, T., Schiele, B.: Learning people detection models from few training samples. In: CVPR. (2011)
- 22 Kaneva, B., Torralba, A., Freeman, W.: Evaluation of image features using a photorealistic virtual world. In: ICCV. (2011)
- 23 W. Li and M. Fritz. Recognizing materials from virtual examples. In European Conference on Computer Vision (ECCV), 2012.
- 24 V Mnih, K Kavukcuoglu, D Silver, A Graves, I Antonoglou, D Wierstra, M Riedmiller. Playing Atari with Deep Reinforcement Learning. 2013. arXiv:1312.5602
- 25 Mnih, V. Kavukcuoglu, K. Silver, D. Rusu, A. Veness, J. Bellemare, M. Graves, A. Riedmiller, M. Fidjeland, A. Ostrovski, G. Petersen, S. Beattie, C. Sadik, A. Antonoglou, I. King, H. Kumaran, D. Wierstra, D. Legg, S. Hassabis, D. Human level control through deep reinforcement learning. *Nature*. 2015; 518, p 529–533.
- 26 Mnih, V. Badia, A. P. Mirza, M. Graves, A. Lillicrap, T. P. Harley, T. Silver, D. Kavukcuoglu, K. Asynchronous Methods for Deep Reinforcement Learning. 2016. arXiv:1602.01783
- 27 Silver, D. Policy Gradient. [Lecture] University College London. 21 Dec 2015. Available at <http://www0.cs.ucl.ac.uk/staff/d.silver/web/Teaching.html>. [Accessed 01/02/2018]