

Intelligent Physiotherapy Through Procedural Content Generation

Shabnam Sadeghi Esfahlani and Tommy Thompson

Department of Computing & Technology
Anglia Ruskin University
Cambridge, UK

shabnam.sadeghi-esfahlani@anglia.ac.uk, tommy@t2thompson.com

Abstract

This paper describes an avenue for artificial and computational intelligence techniques applied within games research to be deployed for purposes of physical therapy. We provide an overview of prototypical research focussed on the application of motion sensor input devices and virtual reality equipment for rehabilitation of motor impairment: an issue typical of patients of traumatic brain injuries. We highlight how advances in procedural content generation and player modelling can stimulate development in this area by improving quality of rehabilitation programmes and measuring patient performance.

Introduction

In recent years computational intelligence (CI) and artificial intelligence (AI) have placed an increasing focus on applications within games. This research has led to a number of sub-fields to emerge, notably procedural content generation (PCG): the algorithmic process of constructing artefacts of some functional or cosmetic value. PCG is often adopted as a generic term that encompasses a variety of procedures whose applications range from the creation of visual aesthetics to the more ambitious realms of automated game development (Togelius et al. 2011). In each instance, there are often regulations or rules imposed upon the generative systems in play: ensuring that content respects some pre-defined expectations. These regulations may be imposed due to the functional needs of the environment the content is crafted for, or due to bias¹ inherent either in the human designers creative process or the computational and mathematical models adopted for content creation (Togelius et al. 2013). These rules have a subsequent impact on the generative space of the final system and the types of content we can build for this problem. In essence, this can ultimately inhibit the spaces within which generation will occur.

It is within this flexibility that we explore an interesting problem space for PCG research in physiotherapy: physical rehabilitation programmes devised to aid patients in improving or regaining motor functions. Physiotherapy is adopted in a wide variety of circumstances, such as personal injury or

disability ranging from a variety of mental or physical complications. In order to maximise the gains of these rehabilitation programmes, patients must be provided with bespoke care and training that is reflective both of the patients current circumstances, but of the goals and expectations of the patient and their therapist.

The last 10 years has seen a significant increase in the use of ‘serious games’¹ for the purposes of therapy and training, be it physical or mental (Halton 2008). In any case, we typically require scenarios within these simulations that cater towards the patients needs. This process is highly reliant on the expert-knowledge of physicians in order to achieve or provide momentum towards a specific goal in mind for the patient. Despite this expertise, there is tremendous risk associated with virtual training scenarios for physiotherapy, given the need to ensure it is scaled to the abilities of the patient - so as not to cause injury - but also with respect to the aspirations and goals of their therapy - thus ensuring they actually improve. On a less critical note, there is the issue of maintaining a patients interest in the process and aiming to further their goals while introducing diversity into the training programme.

With these considerations in mind, this paper highlights prototypical procedural content generation research for physiotherapy; with a specific focus on the rehabilitation of stroke patients. We highlight the problem area as well as existing research in the field of serious games for rehabilitation and the work to-date adopting AI and CI methods to stimulate these processes. This is followed by an overview of the prototypical system devised and our current research in this area. We conclude this paper with a brief discussion of how PCG combined with player modelling research can prove highly influential in games of this nature.

Background & Related Work

A stroke occurs when the blood supply to a particular part of the brain is suddenly cut off, resulting in cells in the affected area of brain become damaged (National Health Service). While potentially fatal, there are over 1.2 mil-

¹A term adopted for games with a primary interest in areas such as education and/or training rather than entertainment. The term is often attributed given the stigma associated with games as a medium.

lion stroke survivors currently residing in the United Kingdom (UK) with approximately 152,000 new cases identified each year (Stroke Association 2016). Stroke survivors suffer from a number of symptoms, but some of the most long-term damaging include reduced mobility and a lack of co-ordination and balance. As a result, many stroke sufferers are the recipient of physiotherapy treatment in an effort to return to pre-stroke levels of mobility. This is a costly and time consuming process estimated to cost the UK economy 9 billion a year (Halton 2008). Regardless of economic factors, the success of this treatment is variable due to the nature of strokes themselves: with some patients requiring longer periods of physiotherapy before regaining their former independence. In addition, the continued motivation of the patient in working towards their goals (National Health Service) is equally important.

In recent years, Virtual Reality (VR) has been identified as having potential therapeutic benefits (Halton 2008). Through use of VR technology, clinicians can provide an environment to practice and monitor repetitive motions and provide feedback to guide the rehabilitation process. This has been found to improve rehabilitation for specific motor-learning tasks of patients suffering from traumatic injuries (Ustinova et al. 2014). This includes the movement of hands, upper and lower limbs (Carlozzi et al. 2013; Saposnik et al. 2010), unilateral and bilateral movement (Mirelman, Bonato, and Deutsch 2009). Given that these platforms can range from high-cost bespoke systems to low-cost off-the-shelf video gaming technologies, it presents an opportunity - albeit with risk - for patients to perform additional exercises at home without the supervision of the therapist (Koenig et al. 2014). Lastly, it allows patients to focus on the training process given the immersive nature of the simulated environment (Dobkin 2004; Holden and Dyar 2002).

In many respects the ideas of rehabilitation-driven games share similarities with that of dynamic difficulty adjustment. However, as noted in the likes of (Burke et al. 2009), the challenges faced are more nuanced given the physical challenges faced by patients. A notable example can be found in the Intelligence Game Engine for Rehabilitation (IGER) detailed in (Pirovano et al. 2012; Borghese et al. 2013) which is focussed on the recording and management of physiotherapy for stroke patients. This system adopts an adaptive fuzzy-driven approach to catering for the patients rehabilitation process. This is achieved while working with a variety of small mini-games that utilise peripherals such as the Microsoft Kinect and Nintendo Wii Fit board. This system places emphasis on ensuring that the parameters of the current game are such that it should not result in activities that player might find painful.

Project Overview

The prototype gameplay system shown in Figure 1 is designed to aid stroke patients to stimulate motor movement. While our long-term efforts seek to adopt the use of virtual reality headsets², our initial testing - and focus of this pa-

²VR headsets are already implemented for this system, but is beyond the scope of current research interests

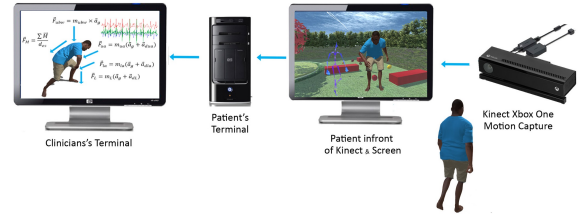


Figure 1: A summary of the current prototype system: in which patients can interact with a game-environment courtesy of the Kinect interfaces.

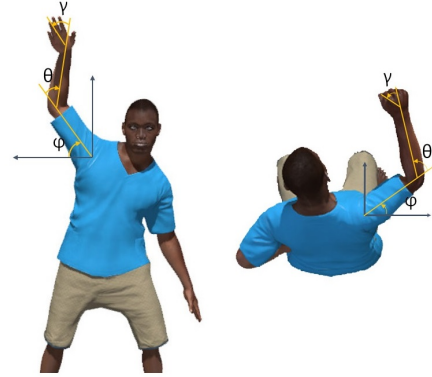


Figure 2: The Kinect and Myo sensors allow for accurate representation of the horizontal and vertical motion of the player.

per - is working solely with gesture recognition apparatus: namely the Microsoft 'Kinect' camera and 'Myo' wearable technologies³. The Kinect camera allows for basic gesture expressions to be captured with Kinect, while the Myo provides more accurate gyroscopic data in addition to muscular activity. This allows for movements to be monitored and co-ordination abnormalities to be recognised through subsequent analysis at a rate of approx. 30 frames per second.

The adoption of Kinect and Myo allows us not only to model the patient within the scene but enables full-body interaction including upper and lower limbs (Figure 2). Through these sensors, the participant is able to manipulate the behaviour of a player (avatar) within the environment as shown in Figure 2. Subsequent therapy sessions are presented using a series of mini-games that require specific kinematic activity such as that shown in Figure 3. At the time of publication, two simple games are implemented. The first is a button-pressing game where players must hover their hand in the location to enable them. The second is a fruit grabbing game (Figure 3) where players must grab objects and drop them in the basket at their feet. These games only requires simple single-limb guided movements. However, future plans aim to adopt complex, whole-body task-oriented actions such as achieving specific posture, balance,

³<http://www.myo.com>



Figure 3: A screenshot of the current working prototype, with an emphasis on interaction with objects in the virtual 3D space.

eye-tracking and facial expression tasks.

With these components in place, a control system has been established for the purposes of conducting physical rehabilitation exercises. However, it is the usage of patient data both in terms of current kinematic and muscular performance and a developed understanding of the needs of the patient could yield a more comprehensive and reactive physiotherapy programme. It is at this juncture we argue that computational intelligence methods found within games research that could prove useful in elevating this work.

Our ultimate goal is for both patients and physicians to interact with the system. While patients continue to play these mini-games and can see their progress against increasingly challenging tasks, physicians can assess the progress on a more granular and data-driven level. This will help identify ranges of motion that players have either not yet fully explored or have not performed well within, allowing for the physician to then tailor the system parameters to move towards more specific goals in-line with their expectations of the patient.

Discussion

The intention of this prototype is to enhance stroke therapy in an engaging and safe fashion. Our aim is to allow participants with different movement capabilities to play the game and at the same time minimize any frustration. Furthermore, we wish to address the possibility of providing patients with real-time feedback to enable correction and adaptation of erroneous behaviour. This presents a rather interesting opportunity for procedural content generation.

The inputs as detailed in Figure 2 translates to a large range of motion. However, this space is in-turn limited by the actual capabilities of the player, assuming they are suffering from limitations in mobility. As such, games previously implemented such as the aforementioned button pressing or fruit grabbing are - at the time of writing - not fit for purpose. No constraints are exhibited upon them beyond a range of possible physical space that a player can interact within. They do not take into consideration the perceived and desired mobility of the player. This leaves us with numerous challenges to face, which we will now briefly summarise in addition to some initial thoughts on tackling them.

Perceived and desired mobility

At minimum there are considerations to be made for generating within a restricted area of the player: understanding the range of comfortable motion within which the player can move should be used to automatically constrain the generative system with respect to the physical placement while minimising the impact upon expressivity of the generation itself (Togelius et al. 2011). There is precedent for the adoption of parametrisation in order for players to dictate generated content, most notably the Mario AI level generation track (Shaker et al. 2011): in which data acquired from a preliminary playthrough is adopted as a basis for subsequent level creations. While these parameters are based on player performance in the case of *Super Mario*, we consider the range of movement exhibited by players in preliminary 'warm-up' games as means to manage future behaviour. The subsequent generation can be tackled through use of vanilla genetic algorithms to isolate potential parameters for object placement, with constraints adopted as heuristics. This could result in a variety of levels for games - or even mini-games themselves - that constrain themselves with respect to the perceived motion of the patient. This would help not only maintain an element of novelty in the terms of player engagement, but maximise the range of motions being adopted by the patient.

Physicians Expectations

There is an interesting opportunity in the adoption of player-driven metrics in this problem space, given that the 'levels' or tasks crafted for a particular play session should be geared towards not just the players current ability, but also the expectations of their physician. Furthermore, there is the additional consideration of ensuring that the games crafted also challenge these actual and perceived limitations in such a fashion that it results in continued gains. One possibility is to measure and curate levels/games based on how well they may suit specific tasks. This is critical for it to prove practical as part of the physiotherapy process, given it would help maintain a critical focus for future therapy sessions.

The Risk of Injury

A crucial aspect of this research is minimising the risk of injury. Unlike most procedural and adaptive play domains, we are dealing with circumstances that are safety-critical in nature and any interaction with the system could prove equally as detrimental to a patient as it could be beneficial. As such, constraints imposed upon the generative system must ensure any risk presented towards players is minimised. This is critical given not only the need to ensure no harm comes to the patient, but to retain their trust in a systems aimed at assisting them in their development. While a more robust means to achieve this is yet to be established, the formative 'practice-run' mentioned earlier would prove useful at establishing just how far a player will move in these sessions: given that patients will typically not deliberately cause themselves injury.

Patients Perceptions of Ability

Having previously discussed the issue of managing physicians expectations, there are also the challenges in addressing the physical and mental limitations of the patient. Players may err on the side of caution for fear of injury, which will subsequently impede their development. There are two issues to consider in this regard: the perception of whether the generated content could prove harmful to the user and whether the user believes the content is beyond their ability. The modelling of such qualitative statements in reaction to the specified metrics has previously been tackled through the likes of fuzzy rules (Pirovano et al. 2012) as well as neuro-evolutionary preference learning (Shaker, Yannakakis, and Togelius 2011). In each case, levels are qualified through qualitative metrics: an approach that can not only establish perceived difficulty, but a similar approach can be adopted for expectations of therapists. By training systems to establish what is feasible versus what is ‘achievable’ in the minds of players, we can isolate problem areas where players may still gain tangible improvement.

Conclusion

This paper highlights early-stage research in the development of a game-driven framework for physical therapy and rehabilitation: highlighting the potential problem areas and challenges faced when developing generative systems that act within a safety-critical system. This work is still in its formative stage, with the initial phase focussed on players providing initial data to aid in configuring subsequent level generation currently implemented, with future work focussed on establishing models of perceived difficulty for generative systems to tailor games for individual players.

References

- Borghese, N. A.; Pirovano, M.; Lanzi, P. L.; Wüest, S.; and de Bruin, E. D. 2013. Computational intelligence and game design for effective at-home stroke rehabilitation. *Games for Health: Research, Development, and Clinical Applications* 2(2):81–88.
- Burke, J. W.; McNeill, M.; Charles, D. K.; Morrow, P. J.; Crosbie, J. H.; and McDonough, S. M. 2009. Optimising engagement for stroke rehabilitation using serious games. *The Visual Computer* 25(12):1085–1099.
- Carlozzi, N. E.; Gade, V.; Rizzo, A. S.; and Tulskey, D. S. 2013. Using virtual reality driving simulators in persons with spinal cord injury: three screen display versus head mounted display. *Disability and Rehabilitation: Assistive Technology* 8(2):176–180.
- Dobkin, B. H. 2004. Strategies for stroke rehabilitation. *The Lancet Neurology* 3(9):528–536.
- Halton, J. 2008. Virtual rehabilitation with video games: A new frontier for occupational therapy. *Occupational Therapy Now* 9(6):12–14.
- Holden, M. K., and Dyar, T. 2002. Virtual environment training: a new tool for neurorehabilitation. *Journal of Neurologic Physical Therapy* 26(2):62–71.
- Koenig, S.; Ardanza, A.; Cortes, C.; De Mauro, A.; and Lange, B. 2014. Introduction to low-cost motion-tracking for virtual rehabilitation. In *Emerging Therapies in Neurorehabilitation*. Springer. 287–303.
- Mirelman, A.; Bonato, P.; and Deutsch, J. E. 2009. Effects of training with a robot-virtual reality system compared with a robot alone on the gait of individuals after stroke. *Stroke* 40(1):169–174.
- National Health Service, howpublished = “http://www.nhs.uk/conditions/stroke/pages/introduction”, y. . . Health and social care information centre (hscic).
- Pirovano, M.; Mainetti, R.; Baud-Bovy, G.; Lanzi, P. L.; and Borghese, N. A. 2012. Self-adaptive games for rehabilitation at home. In *2012 IEEE Conference on Computational Intelligence and Games (CIG)*, 179–186. IEEE.
- Saposnik, G.; Teasell, R.; Mamdani, M.; Hall, J.; McIlroy, W.; Cheung, D.; Thorpe, K. E.; Cohen, L. G.; Bayley, M.; Group, S. O. R. C. S. W.; et al. 2010. Effectiveness of virtual reality using wii gaming technology in stroke rehabilitation a pilot randomized clinical trial and proof of principle. *Stroke* 41(7):1477–1484.
- Shaker, N.; Togelius, J.; Yannakakis, G. N.; Weber, B.; Shimizu, T.; Hashiyama, T.; Sorenson, N.; Pasquier, P.; Mawhorter, P.; Takahashi, G.; et al. 2011. The 2010 mario ai championship: Level generation track. *IEEE Transactions on Computational Intelligence and AI in Games* 3(4):332–347.
- Shaker, N.; Yannakakis, G. N.; and Togelius, J. 2011. Feature analysis for modeling game content quality. In *2011 IEEE Conference on Computational Intelligence and Games (CIG’11)*, 126–133. IEEE.
- Stroke Association. 2016. State of the nation stroke statistics. <https://www.stroke.org.uk>.
- Togelius, J.; Yannakakis, G. N.; Stanley, K. O.; and Browne, C. 2011. Search-based procedural content generation: A taxonomy and survey. *IEEE Transactions on Computational Intelligence and AI in Games* 3(3):172–186.
- Togelius, J.; Champandard, A. J.; Lanzi, P. L.; Mateas, M.; Paiva, A.; Preuss, M.; and Stanley, K. O. 2013. Procedural content generation: Goals, challenges and actionable steps. *Dagstuhl Follow-Ups* 6.
- Ustinova, K.; Perkins, J.; Leonard, W.; and Hausbeck, C. 2014. Virtual reality game-based therapy for treatment of postural and co-ordination abnormalities secondary to tbi: A pilot study. *Brain injury* 28(4):486–495.