

Review Article

Copyright © All rights are reserved by Mc Leod RD

A Framework for Utilizing Serious Games and Machine Learning to Classifying Game Play Towards Detecting Cognitive Impairments

Gutenschwager K¹, Shaski RA², Mc Leod RD^{2*} and Friesen MR²¹Institute of Information Engineering, Ostfalia University of Applied Sciences, Germany²Department of Electrical and Computer Engineering, University of Manitoba, Canada***Corresponding author:** Mc Leod RD, Department of Electrical and Computer Engineering, University of Manitoba, Canada.**Received Date:** November 23, 2019**Published Date:** December 05, 2019

Abstract

Objective: This work presents a framework for geriatricians with interests in exploring the opportunities that serious games and machine learning may offer in assisting with diagnoses of cognitive impairment. Mild Cognitive Impairment (MCI) is often a precursor to more serious dementias, and this work focuses on the potential of applying machine learning (ML) approaches to data from serious games to detect MCI. The focus on MCI detection sets the work apart from mobile games marketed as 'brain training' to maintain mental acuity.

Materials & methods: Our prototypical game is denoted WarCAT. WarCAT captures players' moves during the game to infer processes of strategy recognition, learning, and maintenance (memory, or conversely, loss). As with any ML endeavor, effectively demonstrating this conjecture requires large datasets, and thus, the objective of this work is to also develop ML methods to generate synthetic data that can plausibly emulate a large population of players. To generate synthetic data, the ML paradigm of Reinforcement Learning (RL) is being applied as it most closely emulates the way humans learn. Subsequently, The ML-based classification of game-play data to identify varying degrees of cognitive impairment.

Results: We have developed RL bots that process millions of gameplay training patterns and achieve gameplay results comparable to the best human performance. This baseline allows us to create bots with various levels of training to emulate individuals at various stages of learning, or by extension, various levels of cognitive decline.

Conclusion: The work introduces a framework to demonstrate the characteristics and potential that ML offers to cognitive health diagnosis. Specifically, the work demonstrates the RL work and the ML methods to successfully classify different levels of gameplay.

Keywords: Mild cognitive impairment; Machine learning; Reinforcement learning; Serious games

Introduction

Mild Cognitive Impairment (MCI) is often a precursor to more serious dementias, and early detection is of significant value to therapeutic intervention. Serious games can be a rich source of data for mental health, [1,2] and while there are other widely deployed games on the market that claim to do 'brain training' to maintain cognitive function or potentially delay MCI, no widely deployed game has been developed to *detect or help assess* MCI - a subtle but important distinction. It is well recognized that early detection of any health concern usually leads to a more positive prognosis, and early diagnosis of cognitive impairment is no exception [3,4] This work focuses on the potential of analyzing game-players data toward MCI detection. Although there is considerable research activity in developing and analyzing serious games for cognitive

health, [5-10] the role of serious games for any range of genuine health applications is embryonic at best and considerably ad-hoc, with some effort being made to more formally organize the topic [11].

For these purposes, this paper briefly describes a serious game developed within our research group, denoted War CAT. It is presented only as a platform for the machine learning (ML) work that follows, in that the objective of this paper is to discuss the machine learning approaches used to generate synthetic game-play data using Reinforcement Learning approaches, and initial attempts at ML-based classification of game-play data for various degrees of impairment. The paper outlines the ML approaches used, including the justifications, opportunities, and trade-offs. It

also explores other directions available for ongoing work or work that others with similar interest in the healthcare intersection of SG and ML may undertake.

The present work was preceded by extensive investment into the development of a simple, mobile serious game (WarCAT) suitable for MCI detection [12]. Ultimately, the aim as with many games is to make War CAT widely available and subsequently to build a dataset of human users, from which to detect baseline

cognitive functions and pre-symptomatic MCI using ML approaches on real data. The present step aims to develop these ML approaches, which rely inherently on large datasets. To address this circular dilemma, the present work explores ML approaches to synthetic data generation, to then be used in ML approaches of data classification to emulate various stages of MCI. As the present work does not use any human or animal subjects, no institutional ethics board review was required. The overall framework invalidating an SG ML classification technique is illustrated in Figure 1.

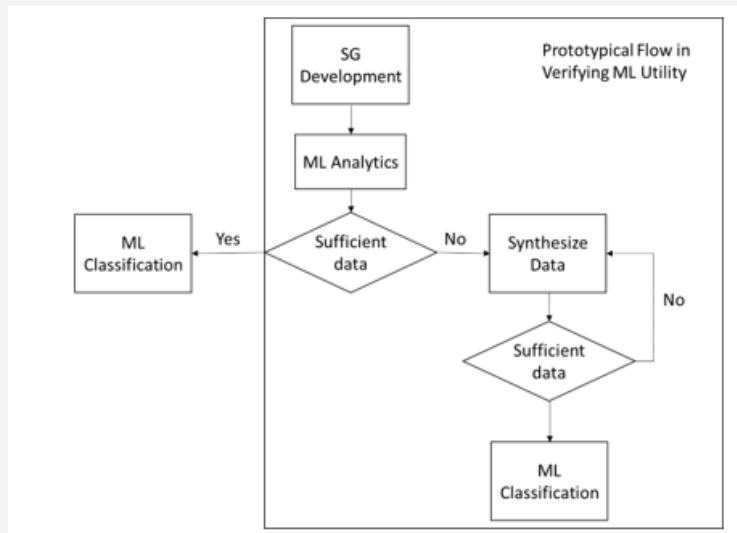


Figure 1 : Framework for verifying Serious Games (SG) and Machine Learning (ML) utility.

WarCAT platform

Serious games for these purposes arguably should not be overly complex. WarCAT is a simple game, a modification of the familiar card game of WAR. Figure 2 illustrates an instance of play during

one round of WarCAT. In WarCAT, the human player is playing with five cards that can be played in the order of their choice against a bot or computer that is playing in a deterministic fashion (i.e. using a set strategy such as low to high cards, high to low cards, etc.).




Figure 2 : An instance of play during WarCAT.

WarCAT illustrates a game's ability to collect data that can, in turn, be analyzed to capture some degree of a player's cognitive function related to strategy recognition, development (learning), memory, recall, and other executive functions. This can be represented in

a data visualization of a person's "cognitive fingerprint," and this visualization provided the motivation to pursue more algorithmic means of classification through ML approaches.

Generating synthetic data

A difficulty faced by all ML efforts in extracting meaningful inferences from game play data is data volume. Towards this end, researchers rely heavily on generating synthetic data to help establish the utility of their ML classification techniques. Of the various ML approaches available, Reinforcement Learning (RL) is ideally suited to generating synthetic data to emulate cognitive performance and decline. Of all ML strategies, RL is the one that most closely emulates how humans learn many tasks. In general, a stimulus is presented, an action taken, and a reward emitted as illustrated in Figure 3.



In the case of WarCAT, the stimulus is the cards played by the computer (game bot) and the action is the card played by the player (player bot). The objective is then to build player bots to emulate players with various levels of cognitive performance and/or cognitive decline. Although RL has a long history and various mathematical manifestations, the principles are simple. Player bots learn through repeated play where positive rewards reinforce actions that emit further positive rewards.

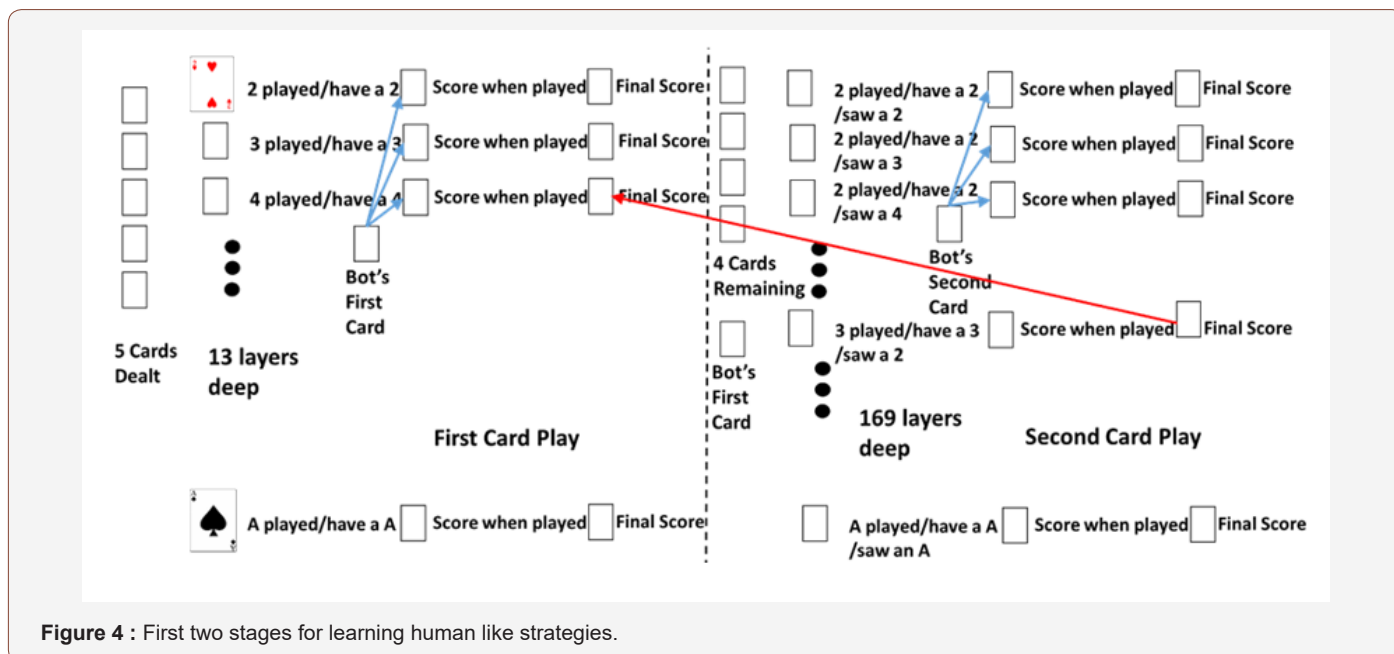
In the traditional card game of WAR, the highest played card would win. A scoring variation is used here to encourage greater

diversity of strategy – that is, the scoring is more heavily weighted to a narrow victory over a larger difference in card values. In this manner, an Ace beating a King is of greater value than an Ace beating a Two. The following equation illustrates the scoring for each card played within a hand:

$$\text{Score} = 13 - (\text{Winning Card Value} - \text{Losing Card Value}), \quad (1)$$

Unless there is a tie, in which case the score for that play is 0. A positive reward is associated with winning, while a negative reward or score with losing. With this scoring in place, it is easier to learn a winning strategy through reinforcement if one wins more often.

The approach takes its lead from Artificial Intelligence (AI) techniques being applied to entertainment game agents or bots. In the case of AI agents in games, typically the objective is not to outsmart players or build super bots. Rather, the objective is to have the agents be as smart as necessary to maintain the human player’s engagement and sense of challenge and fun. The objective is not to push the limits of ML within the bot, as would be the case in many ML optimization instances. For our purposes, the ML bot needs to be imperfect, imitating a human-like behavior. With this objective, the goal is to have bots train for various periods or various number of training patterns or iterations. Bots trained with minimal iterations will essentially have little knowledge of what to do, while those that have been trained longer will have learned strategies that enable them to win consistently. As such, bots with various levels of training will represent human play with various levels of cognitive skills or conversely, cognitive impairment.



An illustration of the proposed Reinforcement Learning method is shown in Figure 4, which represents the more significant combinations of conditional play with the first two cards played. Strategies of play are learned by playing against a game bot that is dealt with five random card hands where the game bot plays its cards high to low. Scoring boxes are shown, which are populated

during training; these include the immediate score when each card is played as well as the total score for that hand upon completion of the hand. Effectively this process resembles that of finding the most suitable paths or trajectories through a large decision network (Figure 5). Effectively, the graph would initially be a complete network of five layers of 13^5 , 13^6 , 13^6 , 13^6 , and 13^5 nodes

respectively, representing the first to last card played. As such, each layer corresponds to the card to be played, i.e. the layers in the middle are not conventional hidden layers as they also have associated inputs as the cards are played. The number of nodes

has been reduced in the figure by a factor of 13^4 for clarity. This simple RL method is one of many that are suitable for training bots to emulate human players.

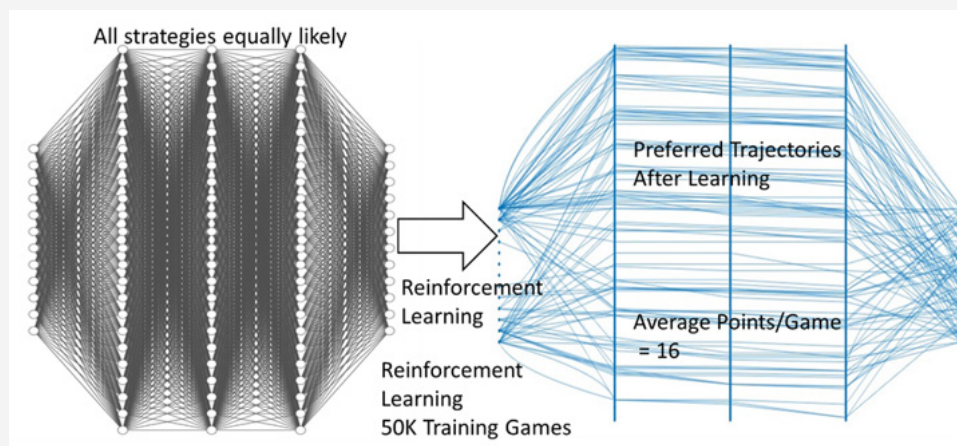


Figure 5: Emergent human like strategies.

With a simple Reinforcement Learning algorithm, preferable game play trajectories emerge. During training, cards are initially played at random and scores tabulated. After a series of games, preferable strategies become more heavily weighted. All moves are conditional on the cards the player bot has been provided as well as the previous card(s) the game bot played. For this initial foray into using Reinforcement Learning, only the last card played by the bot was used as a condition. After enough training, the more favorable trajectories through the game can be used to accelerate the reinforcement. This phase, driven by learned pdfs is analogous to a self-organizing network. The decision on what card the player bot plays could also be a function of the immediate scores accumulated during training. The criteria for learned decision making is not the most obvious in relation to real play by an individual, who may also be more risk averse or risk tolerant. However, it is likely that decisions are made based on the cards held, as well as what the game bot most recently played.

Figure 5 also illustrates the original network of strategies (left-hand figure, in which all sequences are equally likely) and the network of strategies that emerge as being more favorable (right-

hand figure). The clustering in the left-hand side of the trained network (right-hand figure) illustrates the initial sacrificing of a low card or the playing of a high card if held.

Computational results for synthetic data

The first computational results presented are for the following scenario: player bots are to play against game bots, with the game bot playing a consistent strategy of its cards, that being of high to low. Each is dealt with five cards at random from independent decks of 52 cards. The values associated with the cards are 0 through 12.

Although a variety of simulations could be undertaken, most will have similarities to those presented here. The basic Reinforcement Learning strategy is to present the player and game bots with a few hands, whereby the game bot plays high to low and the player bot plays randomly. This period of play is denoted an epoch, followed by another epoch of play where the player bot uses what is learned from the first epoch to adapt further to improve its strategy. This process is simple enough to provide insight into any performance advantage associated with adaptation. The number of training patterns range within an epoch range from 10 to 10^{11} .




Figure 6: Average Score per Player Bot over 150 Hands.

Training is followed by a testing phase. Tests consist of a number of player bots playing games of 150 hands, similar to the number of hands played by real players in the current manifestation of the game. The number of player bots has been set at 100 for the testing phase. Decisions made on which card to play by the player bot is simply the card that is associated with the greatest reward as opposed to making a probabilistic decision weighted by the reward associated with each card.

Figure 6 illustrates the average score per player bot when playing against a game bot for 150 hands. As expected, the player bot improves considerably as the number of training patterns increase. It may also be argued that there is diminishing return after 10^{10} training patterns as the network may be beginning to memorize previously encountered patterns. Memorization is a common theme or difficulty associated with ML and is important to recognize.

Table 1 illustrates the average scores and variance for 100 player bots achieved over 150 hands of play as a function of the degree of RL or training patterns. As the game is stochastic there is considerable variation in the overall average score. This variation will present a challenge for any machine learning classifier. The variation can be somewhat mitigated by normalizing play against the best score possible, which is, on average, calculated at a value of 25, again with some variation associated with the hands dealt. Figure 7 illustrates an example of normalizing the scoring against the maximum possible per hand for each player.

Table 1: Average score and variance for 100 player bots over 150 hands at various degrees of RL.

Training Patterns	Average Score	Variance
20	0.032	6.405
200	0.338	7.554
2.00E+03	2.461	7.545
2.00E+04	10.695	3.404
2.00E+05	13.255	1.516
2.00E+06	16.343	1.649
2.00E+07	18.921	2.333
2.00E+08	19.704	2.124
2.00E+09	19.273	1.996
2.00E+10	19.332	2.177
2.00E+11	17.762	2.14

The reduction in variance seen in Figure 7 is a result of the difference between two Gaussian distributions (one associated with maximum scores possible given the cards that were dealt and the score the player bot achieved after learning). The variance is reduced as there is a high covariance between the best score obtainable and the score the player bots achieved. This observation lends additional credibility to an ML classifier as the classification of players with various levels of training should form more recognizable clusters with less overlap. Average scores obtained by players is an important statistic to retain as a feature for clustering, and reducing variance is a means of providing better discrimination capabilities for either classical or ML classifiers.

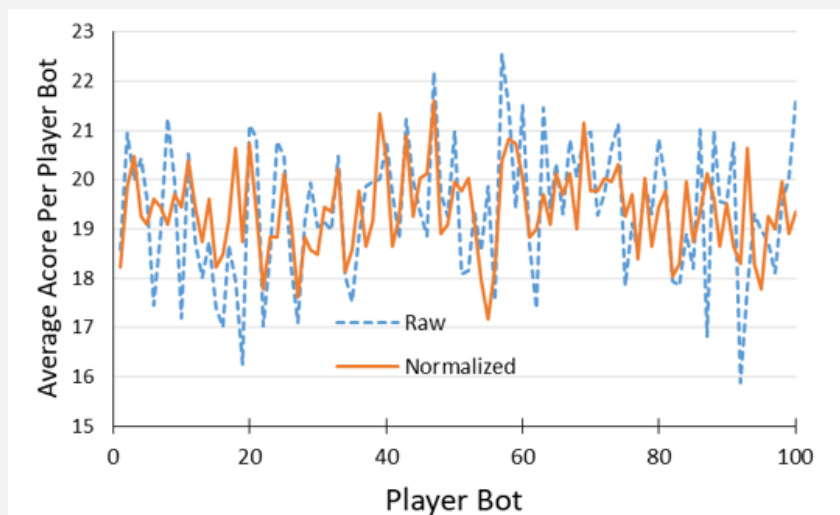


Figure 7: Reduction in variance among players (2×10^{11} training patterns).

Classifying Synthetic Data

The actual ML classifiers must work with hand-by-hand data for real players as well as player bots. An example of hand by hand data and its preprocessing is illustrated in Figure 8.

We have denoted these patterns of play as “cognitive fingerprints”, representative of patterns of play associated with various levels of learning. The conjecture is that the differences in these “cognitive fingerprints” will be significantly easier to classify using the normalized score per hand than the actual scores. An

interesting artifact of play is that although the scores appear highly erratic (Figure 9), they are significantly quantized when compared to the maximum score achievable for that hand. The reason for this effect is that the quantization is directly related to patterns of play associated with misplayed cards. These represent permutations of - at most - three misplayed cards, which results in score differences of 13, 26 or 39 points. These indicate that the play was as learned, but the cards dealt (stochastically) resulted in the deficit. These fingerprints become the features that are used for ML classification.

Its theory it may be possible for a bigger or deeper or more highly tuned classifier to learn this variance reduction and normalization on its own. But as it is an easy data preprocessing step, this type of

data manipulation is recommended as it lessens the computational burden otherwise placed on the classifier.

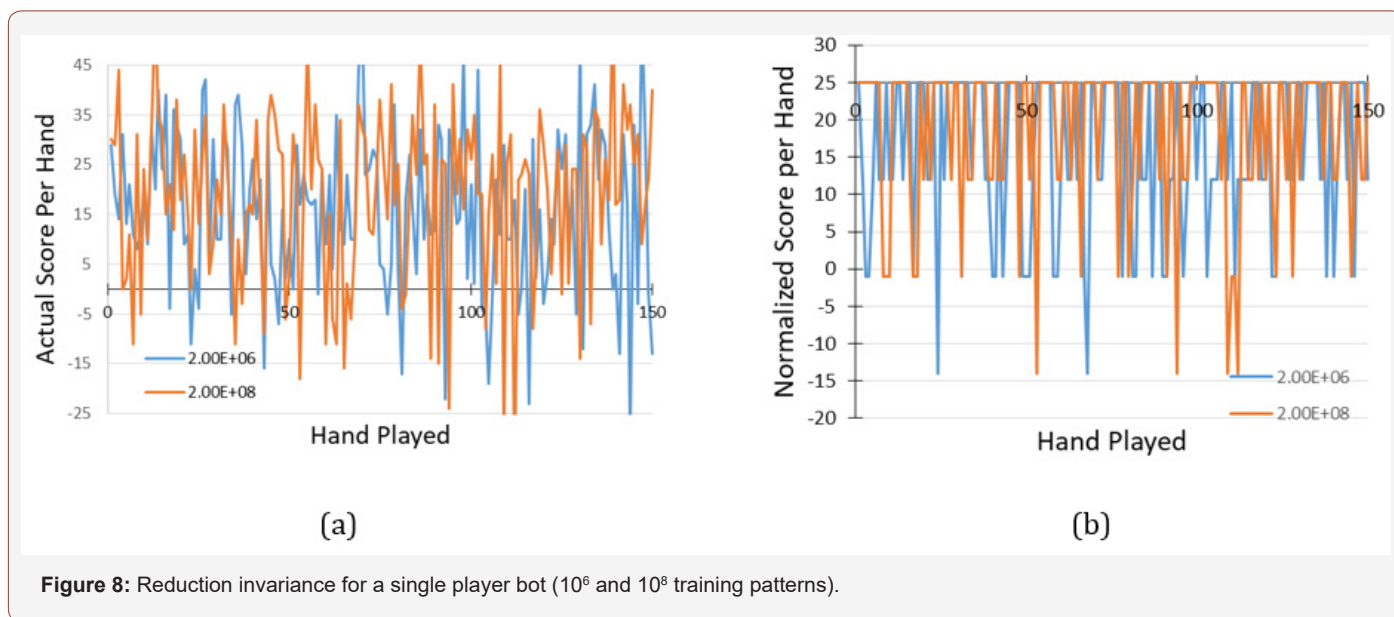


Figure 8: Reduction invariance for a single player bot (10^6 and 10^8 training patterns).

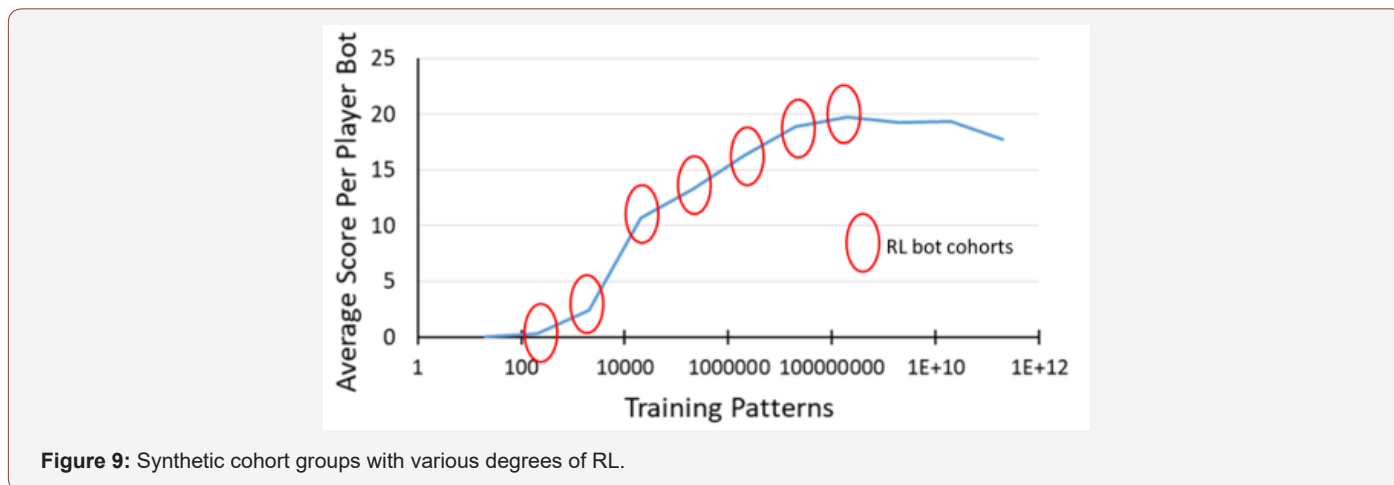


Figure 9: Synthetic cohort groups with various degrees of RL.

Although the promise of ML techniques for classification is enormous, there is an onus on the implementer to process the data in such a manner that the classifier can more easily accommodate. One of the classifier types identified for use here is a convolutional neural network (CNN) [13]. CNNs are traditionally used for image classification but also have corresponding one-dimensional variants [14]. CNNs are supervised learning classifiers based on classifying patterns based on labeled training data.

Classification of the Synthetic Data

For the model here, several cohort's groups were generated based upon their degree of RL. This is pictorially illustrated in Figure 9. Within each cohort group player bots have fingerprints of play like those of Figure 8.

In order to assign a degree of face validity to the bots' play, a visualization was introduced that would break downplay into levels of play. Figure 10 illustrates a visualization developed to help facilitate the human interpretation of how well a person played.

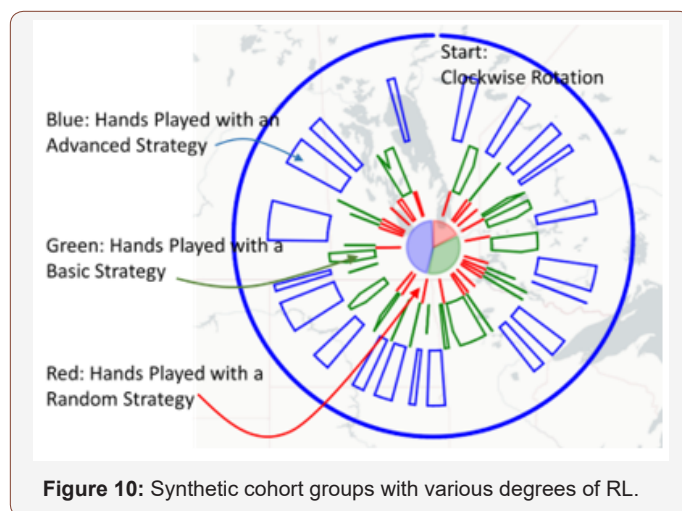


Figure 10: Synthetic cohort groups with various degrees of RL.

Data are displayed on a polar plot with coordinates (angle) governed by the time of the hand and an inference of the strategy of play. Figure 10 illustrates play by the winner of a for-fun tournament. In this case, most of the play was inferred to be Advanced (blue),

with only a small portion inferred as Random (red) or Basic (green), with Basic play more prevalent in the first half of the game than the second half, and conversely, periods of Advanced play increasing as the game proceeded.

A similar visualization for a player bot and best human play to date is shown in Figure 11. In terms of face validity, there appears

to be some degree of similarity in the visualizations derived from a human player and a player bot. In this case, the player bot slightly exceeded the human player. As noted previously, this is often a goal for AI, where the objective is to have better-than-human performance. In our case, the objective is rather to produce artificial behavior emulating humans, and in this case with interpreting cognition with various levels of impairment.

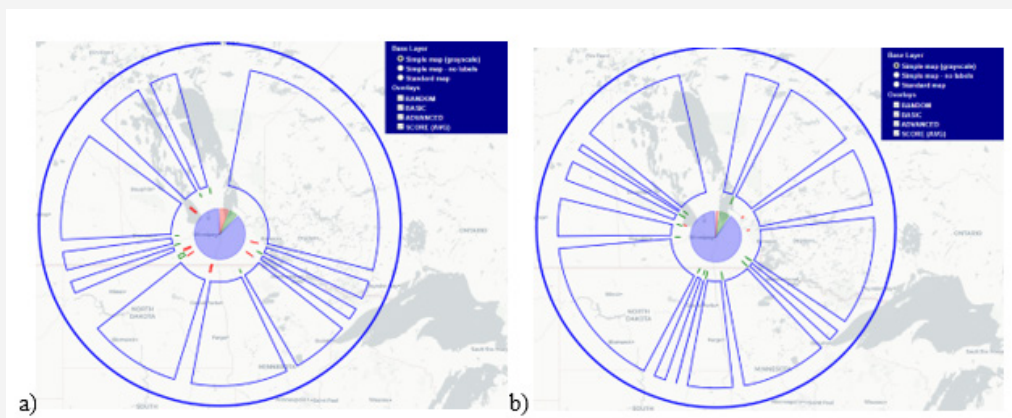


Figure 11: Visualization of (a) well-trained (28 training patterns) bot playing WarCAT, and (b) best human play to date..

The player bot’s “cognitive fingerprint” from Figure 8 combined with aggregate statistics such as average score is suitable features for ML classifiers as the fingerprints resemble spectrograms. Several of the more common approaches in spectrogram classification involve CNNs [15], Recurrent Neural Networks (RNNs) or combinations thereof [16].

As such, the next phase of the research presents and evaluates CNN classifiers for “cognitive fingerprint” images or spectrograms. The CNNs investigated were two convolutional layers, two pooling layers, then five dense layers of different sizes. This is illustrated in Figure 12. These networks were implemented in Python using Google’s TensorFlow high level APIs. The number of synthetic players was 15,000 per cohort. Training and testing were based on 80/20 training/testing partition. A Dense Neural Network (DNN) was also evaluated for completeness.

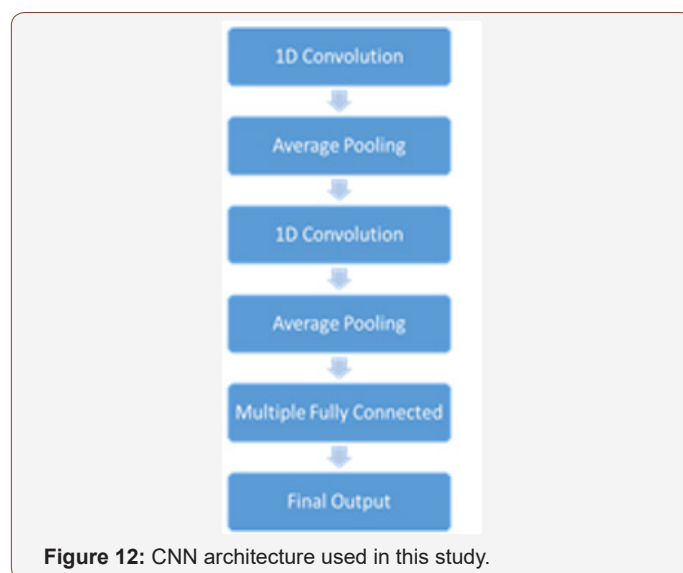


Figure 12: CNN architecture used in this study.

Table 2: Classification accuracy of various CNN architectures and training epochs.

Name	Type	Layers	Epochs	Training Accuracy (%)	Test Accuracy (%)	+/-1 Test Accuracy (%)	Comments
CNN_150	CNN	300,200,	25	97.5	88.9	100	Features: the 150 vectors of residues
		300,200,50,					
		50,36,6					
DNN_150	DNN	216x6,36,6	100	91	84	99.99	
CNN_150_6_35	CNN	300,200,	35	99	88	100	
		300,200,50,					
		50,36,6					
CNN-150_6_50	CNN	300,200,300,	50	99.5	85.8	99.88	Possibly over-trained
		200,50,					
		50,36,6					

Table 2 illustrates various CNN configurations for classifying the cohort bot data generated at various RL degrees of learning, the

training accuracy and the testing accuracy. The +/- 1 accuracy is the accuracy if the misclassification is within one cohort group. This

type of ML classification is reasonably typical with accuracies in the mid-80s. The +/- test accuracy is beneficial as well as expected, and indicates that even if misclassified, it is classified to the next closest cohort group. In most ML initiatives and with the present state of ML technology, considerable parameter and architecture tuning is often required.

Conclusion

This paper provides information to geriatricians who may be interested in the future role that serious games and machine learning may play in assisting in diagnosing cognitive abilities and potential impairment. In most if not all cases, the machine learning methodology will be developed using synthetic data instead of insufficient quantities of real data. Reinforcement Learning is an example among of several techniques for generating data. In this work, the machine learning – serious games framework was demonstrated via a simple mobile game denoted WarCAT. The long-term objective would be to use these types of serious games coupled with machine learning classification to help detect subtle cognitive changes or impairment. The approaches applied to game-play data are akin to building monitoring tools that attempt to recognize subtleties before they are obvious to the ‘person’ or others around them.

Presently, there is no serious game that can detect MCI or any pre-dementia stage of cognitive decline. In any context, if ML is to be effective, considerable data are required. Many areas in which ML techniques have demonstrated utility are those where data are already available, or if not, where data can be generated synthetically. Ideally, one would like to have labeled data such that a supervised learning ML algorithm can be deployed. Of the more effective ML classifiers are CNNs, where they have been incredibly successful for image, music and speech processing. The cognitive fingerprints – based on either real data or synthetic data - have a very close similarity to one-dimensional images or spectrograms and as such may bear similar fruit in ML classification by CNNs. Another strong machine learning contender for this type of data is RNNs as well as hybrids of both.

Although serious games and the challenges of data collection and processing were presented here, it certainly is not the only platform that could be employed to help detect MCI. Many games are likely more suitable, but all will face similar challenges. Other technologies, perhaps more invasive and possibly nefarious, may be built upon Google’s, Amazon’s, or Apple’s assistants. It should also be noted that RL is only one of several means to generate synthetic data; RL just happens to be more widely used in games than alternative techniques.

We also fully recognize that if computer engineers and scientists are to be useful in helping to build technologies for mental health concerns, then collaboration, leadership and guidance from the practicing clinical community is an absolute requirement.

Acknowledgement

Authors RDM and MRF acknowledge the support of the support of Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grants Program and Chairs in Design Engineering Program. WarCAT was developed by Kyle Leduc-McNiven from our group.

Conflict of Interest

No conflict of interest.

References

1. Mc Callum S (2012) Gamification and serious games for personalized health. *pHealth 2012: Proceedings of the 9th International Conference on Wearable Micro and Nano Technologies for Personalized Health* 177: 85-96.
2. Muscio C, Tiraboschi P, Guerra UP, Defanti CA, Frisoni GB (2015) Clinical trial design of serious gaming in mild cognitive impairment. *Front Aging Neurosci* 7: 26.
3. Spaan PEJ (2016) Episodic and semantic memory impairments in (very) early Alzheimer’s disease: The diagnostic accuracy of paired-associate learning formats. *Cogent Psychology* 3(1): 1125076.
4. Petersen RC, Smith GE, Waring WC, Ivnik RJ, Tangalos EG, et al. (1999) Mild cognitive impairment: clinical characterization and outcome. *Archives of Neurology* 56(3): 303-308.
5. Tong T, Chignell M, Tierney MC, Lee J (2016) A serious game for clinical assessment of cognitive status: Validation study. *JMIR Serious Games* 4(1): e7.
6. Hagler S, Jimison HB, Pavel M (2014) Assessing executive function using a computer game: computational modeling of cognitive processes. *IEEE Journal of Biomedical and Health Informatics* 18(4): 1442-1445.
7. Jimison H, Pavel M (2006) Embedded assessment algorithms within home-based cognitive computer game exercises for elders. *EMBS* 06: 6101-6104.
8. Basak C, Voss MW, Erickson KI, Boot WR, Kramer AF (2011) Regional differences in brain volume predict the acquisition of skill in a complex real-time strategy videogame. *Brain Cogn* 76(3): 407-414.
9. Vallejo V, Wyss P, Rampa L, Mitache AV, Müri RM, et al. (2017) Evaluation of a novel Serious Game based assessment tool for patients with Alzheimer’s disease. *PloS One* 12(5): e0175999.
10. Thompson O, Barrett S, Petterson C, David Craig (2012) Examining the neurocognitive validity of commercially available, smartphone-based puzzle games. *Psychology*; 3(07): 525.
11. Verschueren S, Buffel C, Vander Stichele G (2019) Developing theory-driven, evidence-based serious games for health: Framework based on research community insights. *JMIR Serious Games* 7(2): e11565.
12. Leduc-Mc Niven K, White B, Zheng H, McLeod RD, Friesen MR (2018) Serious games to assess mild cognitive impairment: ‘the game is the assessment’. *Research Review Insights* 2(1).
13. Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. *NIPS’12 Proceedings of the 25th International Conference on Neural Information Processing Systems* 1: 1097-1105.
14. Kiranyaz S, Ince T, Gabbouj M (2016) Real-time patient-specific ECG classification by 1-D convolutional neural networks. *IEEE Transactions on Biomedical Engineering* 63(3): 664-675.
15. Costa YM, Oliveira LS, Silla CN (2017) An evaluation of convolutional neural networks for music classification using spectrograms. *Applied Soft Computing* 52: 28-38.
16. Choi K, Fazekas G, Sandler M, Cho k (2017) Convolutional recurrent neural networks for music classification. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* 2392-2396.