# Deep Reinforcement Learning in Immersive Virtual Reality Exergame for Agent Movement Guidance

**Conference Paper** · September 2020

DOI: 10.1109/SeGAH49190.2020.9201901



Aviv Elor
University of California, Santa Cruz

# Deep Reinforcement Learning in Immersive Virtual Reality Exergame for Agent Movement Guidance

Aviv Elor
*Department of Computational Media*
*University of California, Santa Cruz*
Santa Cruz, CA, USA
aelor@ucsc.edu

Sri Kurniawan
*Department of Computational Media*
*University of California, Santa Cruz*
Santa Cruz, CA, USA
skurina@ucsc.edu

***Abstract*—Immersive Virtual Reality applied to exercise games has a unique potential to both guide and motivate users in performing physical exercise. Advances in modern machine learning open up new opportunities for more significant intelligence in such games. To this end, we investigate the following research question: What if we could train a virtual robot arm to guide us through physical exercises, compete with us, and test out various double-jointed movements? This paper presents a new game mechanic driven by artificial intelligence to visually assist users in their movements through the Unity Game Engine, Unity Ml-Agents, and the HTC Vive Head-Mounted Display. We discuss how deep reinforcement learning through Proximal Policy Optimization and Generative Adversarial Imitation Learning can be applied to complete physical exercises from the same immersive virtual reality game. We examine our mechanics with four users through protecting a virtual butterfly with an agent that visually helps users as a cooperative "ghost arm" and an independent competitor. Our results suggest that deep learning agents are effective at learning game exercises and may provide unique insights for users.**

***Index Terms*—Exercise Games (Exergames), Serious Games, Head Mounted Display (HMD), Immersive Virtual Reality (iVR), Project Butterfly (PBF), Machine Learning, Deep Reinforcement Learning, Imitation Learning, Artificial Intelligence**

## I. Introduction

Physical activity is an essential part of daily living, yet 48.3% of the 40 million older adults in the United States are classified as inactive [1], [2]. Inactivity leads to a decline of health with signification motor degradation: a loss of coordination, movement speed, gait, balance, muscle mass, and cognition [1]–[3]. The medical benefits of regular physical activity include weight loss and reduction in the risk of heart disease and certain cancers [4]. However, compliance in performing regular physical activity often lacks due to high costs, lack of motivation, lack of accessibility, and low education [2]. As a result, exercise is often perceived as a chore rather than a fun activity.



Immersive Virtual Reality (iVR) and the increasingly recent use of games for health and well-being have shown great promise in addressing these issues. The ability to create stimulating and re-configurable virtual worlds has been shown to improve exercise compliance, accessibility, and performance analysis [5]–[7]. Other studies have suggested that engaging in a virtual environment during treatment can distract from pain and discomfort while motivating the user to achieve their personal goals [8], [9]. Additional success has been reported in using virtual environments for a broad range of health interventions from a psychological and a physiological perspective [10], [11]. Some of the biggest challenges that these studies found were technological constraints such as cost, inaccurate motion capture, non-user friendly systems, and a lack of accessibility [6], [12], [13].

The past five years have seen explosive growth of iVR systems, stemming from a projected 200 million head-mounted displays systems sold on the consumer market since 2016 [14]. This mass adoption has been in part due to a decrease in hardware cost and a corresponding increase in usability. From these observations, we argue that the integration of iVR as a serious game for health can offer a cost-effective and more computationally adept option for exercise. These systems provide a method for conveying 6-DoF information (position and rotation), while also learning from user behavior and movement. While there has been a number of works in exploring iVR environments for physical exercise [5], [7], [11], we present our paper as an exploration of making these environments more physically intelligent through machine learning. Specifically, we leverage the integration of the Unity Game Engine, ML-Agents, Deep Reinforcement Learning, and a custom in-house iVR exercise game. Through these technologies, we examine how neural network agents can augment a playable experience where a virtual robot arm assists user exercise masked as a task of protecting butterflies from incoming projectiles.

### A. Virtual Reality and Machine Learning

Virtual games provide controlled environments and simulations for a wide range of Artificial Intelligence and Machine Learning applications. Game AI has been extensively researched from mechanical control, behavior learning, player modeling, procedural content, and assisted gameplay [15]. Applying machine learning to the virtual game domain opens up a playground for researchers to find appropriate learning techniques and solve various reward-based tasks [16]. For example, Conde et al showcased reinforcement learning for behavioral animation of autonomous virtual agents in a town [17]. Huang et al demonstrated imitation learning through a 2D GUI to control a Matlab simulated robot in sorting objects [18]. Yeh et al explored Microsoft Kinect exercise with a Support Vector Machine (SVM) classifier for quantified balance performance [19]. Additionally, agent learning in an iVR environment may be especially advantageous for assistive applications.

The computational requirements and data-throughput of modern iVR systems can be leveraged to analyze therapeutic gamification [7], [20], [21], postural analysis [22], and accuracy for research data collections [23]. This is important because iVR systems must have accurate motion capture and low latency of a user's position and rotation from the physical world to reduce motion sickness [24]. As a result, iVR systems are becoming more powerful, immersive, accurate at capturing user behavior, and affordable to the average consumer [14].

Some researchers are recognizing the potential of utilizing machine learning and AI with iVR systems. Zhang et al explored an iVR environment for human demonstrated robot skill acquisition [25]. The authors describe a deep neural network policy to solve this problem for training teleoperation robotics and illustrate that mapping policies of learning using VR HMDs is challenging. Through utilizing an HTC Vive, PR2 Telepresence Robot, and a Primesense 3d camera, the authors successfully trained their neural network to control a robot by collecting user 6-DoF pose and color depth images of player movement. In terms of utilizing machine learning to support player movement, we found two recent studies through our literature review. Kastanis et al described a method of reinforcement learning for training virtual characters to guide participants to a location in an iVR environment [26]. The authors used presence theory to predict uncomfortable interpersonal distance for human players and successfully incentivized study participants to move away from trained virtual agents. And Rovira et al examined how reinforcement learning could be used to guide user movement in iVR through projecting a 6-DoF predictive path for user collision avoidance [27].

While several works have been explored in utilizing machine learning for games, and researchers have started looking at iVR as a medium for human-agent learning, there have been few works exploring agents for iVR exergaming. iVR exercises can provide a vehicle for real-time motion capture and inverse kinematics of player movement. Such data could enable the analysis of confounding postural issues, such as slouched backs and other movement biases, and could adapt the game in real-time to maximize exercise outcome. With these previous works in mind, we consider the following question: what if we could have a predictive model that could inform us of our movement trajectory in a virtual exercise game?

### *B. Study Goals and Contribution*

The prior work discussed in this section has demonstrated that deep reinforcement learning can enable promising predictive models for system control and user behavior. Little work has been done in exploring machine learning from 6-DoF user exercise movement (or movement in general) for iVR experiences. Through this project, which we call "Illumination Butterfly (IB)," we aim to explore how deep reinforcement learning can inform iVR exergames in terms of user movements and game mechanics. Specifically, the goals of this study are to:

1) Examine Deep Reinforcement Learning for a Double-Jointed Virtual Arm to model physical exercise movements through 6-DoF interaction with Immersive Virtual Environments.
2) Explore the capabilities of Generative Adversarial Imitation Learning (GAIL) and Proximal Policy Optimization (PPO) for learning in-game physical exercises.
3) Evaluate the trained agent for cooperative and competitive exercise applications between human users.

Our serious game explores neural network-driven 3DUI interaction techniques by using two emergent machine learning algorithms (GAIL and PPO) to see how a virtual robot arm can both cooperatively and competitively guide users in their movements. This project stems from previous iVR games designed through the interpretation of exercise theory and human anatomy. We expand our work from Elor et al's previous exploration into serious games for upper-extremity exercise movement: a multi-year interdisciplinary exploration between local healthcare professionals, roboticists, game developers, and disability learning centers at Santa Cruz, California [7], [28]–[31]. Through leveraging machine learning, we hope to enable Project IB as a new computational experience to understand human exercise and robotic behavior via virtual butterfly. This project may be a step forward for other researchers interested in integrating "physical intelligence" via predictive models of user movement for other iVR exergames.

## II. System Design

The system in this paper is based on "Project Butterfly" (PBF), a serious iVR game for exercise previously explored by Elor et al [28]. We heavily modified PBF to create a new gaming experience directed at AI guided upper extremity exercises. Our version of PBF was developed in the Unity 2019.2.18f1 Game Engine with SteamVR 2.0 and incorporates the HTC Vive Pro 2018 by Valve Corporation, a highly adopted commercial VR system that uses outside-in tracking through a constellation of "lighthouse" laser systems for pose collection in a 3D 4x4m space [14], [32], [33]. Vive has been verified in previous studies to analyze therapeutic gamification [7], [20], [21], postural analysis [22], and accuracy for research data collections [23].

The objective of the game is to protect a virtual butterfly from inclement weather and projectiles by covering the avatar with a translucent "bubble shield" using the HTC Vive

Controller. Thus the player is required to follow the path of the butterfly with plus or minus 0.1 meters, which enables the dynamic control of pace and position for a prescribed exercise. The player is awarded a score point for every half second they successfully protect the butterfly, with both audio and haptic feedback to notify them that they were successful. By protecting the butterfly, the world around them changes - meadows become brighter, trees grow, and the rain slows down. Conversely, if the butterfly is not protected, no positive feedback occurs - the world does not change. The game can be tailored to each player's speed and range of motion through a dynamic evaluator interface. Previously, PBF was explored with post-stroke and older users to analyze the feasibility of the game with exo-skeletal assistance for two exercises [28] by Elor et al, but was not designed or tested for neural network guided upper extremity movements varying custom exercise movements as reported in this paper.

To explore the application of deep-learning agents for visually guided upper-limb exercise, we created a new modified version of PBF, which included the following changes from the previous version:

1) A modified "Reacher Agent," a double-jointed arm controlled by predictive torque [34], was added into the player controller with the reward given when protecting a virtual butterfly.
2) A training scene for 16 parallel agents and three butterfly movements was created, as shown in Figure 1.
3) A "ghost arm" game mechanic was added for user visual guided movements with the original PBF game modes, and a "human vs agent" game mode was added for competitive analysis.

To the best of our knowledge, this study is one of the first to leverage an immersive VR HMD such as the HTC Vive with deep reinforcement learning to examine visually assisting agents for exergaming.

### A. Machine Learning Environment and Agent Design

Project IB has been fully integrated with Unity ML-Agents, an open-source Unity plugin that enables games and simulations to serve as environments for training intelligent agents. The experimental plugin enables a python server to train agents in development environments through reinforcement learning, imitation learning, neuroevolution, and other emerging Tensorflow based algorithms [32], [35], [36]. We targeted upper-extremity torque and angular momentum as metrics to predict for our model. Having our AI model examine these metrics at the elbow and shoulder joints is advantageous. Torque is important as it used to describe the movement and force produced by the muscles surrounding the joint [37]–[40]. Prior research has examined the torque of upper-body exercise for more in-depth injury assessment; for example, Perrin et al demonstrated that bilateral torque enables clinicians to more accurately set guidelines in the rehabilitation of varying athletic groups [41]. Additionally, angular momentum provides a metric to monitor user movement performance over several exercises, ensuring safety and preventing overuse [42]. Several

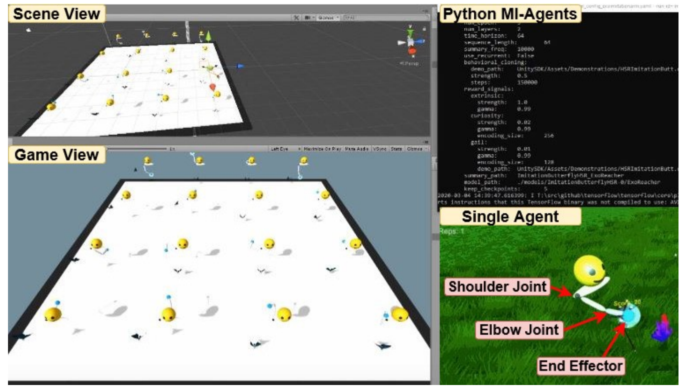


Fig. 1. Project IB Training Scene and AI Agents. Agents act as a double-jointed virtual arm with observation on the shoulder, elbow, and end effector joints. Sixteen agents were set up in parallel to train through the python ml-agents library with an action space of +/- 1.0 for actuating pitch and roll torques on the elbow and shoulder joints, respectively. A reward of +0.01 is given to the agent per every frame the end effector successfully remains on the butterfly. The training scene tasks agents to collectively learn three exercise movements: Horizontal Shoulder Rotation, Forward Arm Raise, and Side Arm Raise.

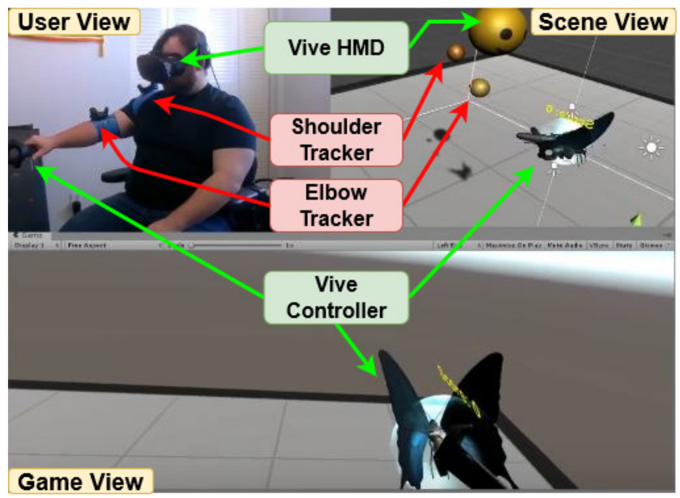


Fig. 2. Project IB Imitation Learning and User Demonstration. A user demonstrates how to protect a butterfly. Vive Trackers are placed on the user's shoulder and elbow joints to record fixed joint movement dynamics. The agent is set to heuristic control to observe the user's joint torques, angular momentum, and hand (bubble) position. A reward of +0.01 is given to the user per every frame the bubble successfully remains on the butterfly. The recorded demonstration is then used to augment reward during parallel agent training with GAIL & PPO.

other studies have explored the benefits of quantifying angular momentum for robotic assistance [43], the severity of lower body gait impairment [44], [45], and how it contributes to whole-body muscle movement [46]. Predicting average torque and angular momentum through an AI model may hopefully provide insights for user movements and future assistive robotic design for Project Butterfly to be re-evaluated with exo-skeletal assistance [28], [47].

With our target predictions in mind, we chose to utilize the Unity Ml-Agents Reacher Agent and Deep Deterministic Continuous Control as it observes and predicts agent fixed

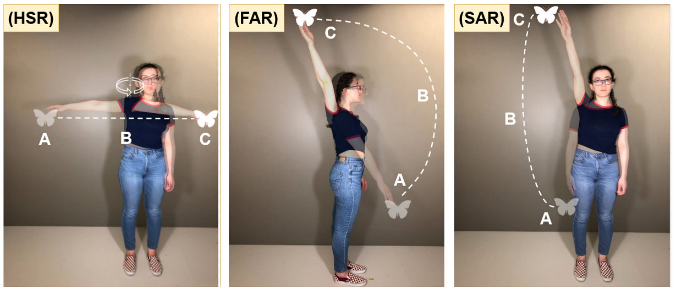


Fig. 3. Project IB exercise movements for Horizontal Shoulder Rotation (HSR), Forward Arm Raise (FAR), and Side Arm Raise (SAR). Movement directions are indicated by the labels ABC followed by CBA for one repetition.

joint dynamics to complete a given virtual task [35], [36]. We modified the agent to act as a double-jointed virtual arm with specific control and observation on the shoulder, elbow, and end effector joints. This allows our agents to collectively learn from an action space from +/- 1.0 where the agent observes joint torques, angular momentum, and butterfly position to predict shoulder and elbow torque. The agent was given a +0.01 reward per every game engine frame update that the bubble or end effector was successfully on the butterfly. Three exercises were targeted for the agent to learn from Horizontal Shoulder Rotation (HSR), Forward Arm Raise (FAR), and Side Arm Raise (SAR), as shown in Figure 3. These movements were chosen as they are considered conventional movement modalities required for active daily living [28], [47].

To examine agent learning, we chose to explore two learning algorithms: Proximal Policy Optimization (PPO) and Generative Adversarial Imitation Learning (GAIL). PPO is a policy gradient method of reinforcement learning that allows sampling parallel agent interaction with an environment and optimizing the agents objective through stochastic gradient descent [48]. GAIL is an imitation learning method where inverse reinforcement learning is applied to augment the policy reward signal through a recorded expert demonstration [49]. In short, GAIL provides a medium for the agent to imitate the user's exercise, and PPO helps the agent find the maximal reward policy to protect the butterfly.

### *B. Agent Training*

Two training sessions were examined through Project IB: parallel agent training (as shown in Figure 1) with PPO only, and PPO with GAIL. We examined the PPO only model to determine the agent performance when solving for maximal reward and the GAIL + PPO model to see if user demonstrations can influence the training process and or personalize agents to the user's movement biases. For GAIL, a demonstration was recorded for each butterfly exercise movement by a human demonstrator, as shown in Figure 2. To record human demonstration, a user was tasked with demonstrating to the agent how to protect the butterfly through arm movement. Vive Trackers were placed at the user's elbow and shoulder joints for agent observation of movement dynamics. This was achieved by creating virtual fixed joints in Unity and inputting

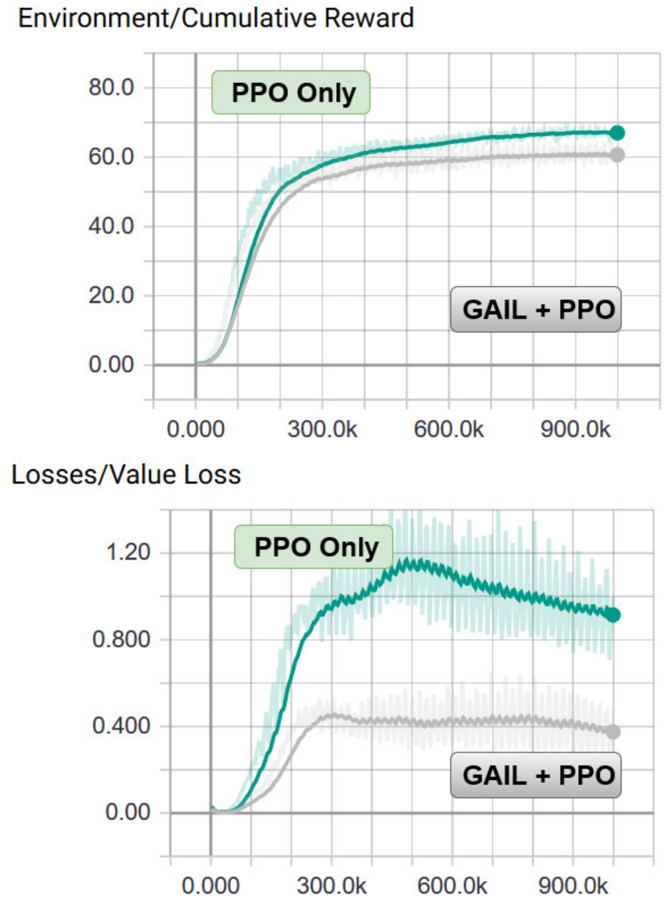


Fig. 4. Project IB Training Results from Tensorboard for one million steps. Results are viewed from the cumulative 16 agents trained in parallel for the three PBF exercises. The "PPO Only" model attained the highest reward with a 11.4% increase compared the "GAIL + PPO" model. Darker lines indicate smoothed results and lighter lines indicate raw data.

rigid body torque and angular momentum into the heuristic agent model. Users demonstrated ideal movements to the agent for about two minutes per exercise.

Training was done with sixteen agents in parallel, as shown in Figure 1. Model parameters were tuned to each trainer_config.yaml file as recommended in the Unity ML-Agents v3.X.X plugin [35], [36]. The training parameters differed between "PPO Only" and 'GAIL + PPO," where GAIL was added as a parameter to the PPO reward signal with a strength of 1%. Full tuning parameters and trained models can be found at *https://github.com/avivelor/UnityMachineLearningForProjectButterfly*. Each training model was run for one million steps at a time scale of 100 through the unity ml-agents API. This was equivalent to about a couple hours of training per each model where agents attempted to learn Horizontal Shoulder Rotation, Forward Arm Raise, and Side Arm Raise.

### *C. Training Results*

Training results between the two models can be seen in Figure 4. Both models demonstrated a promising learning

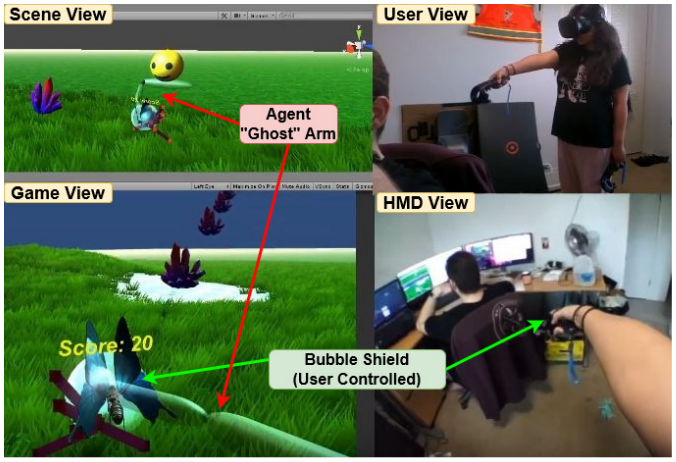


Fig. 5. Project IB Cooperative Gameplay with Trained Agent. The user controls the bubble shield through the controller as a transparent "ghost" arm appears through the user to help guide and predict user movement in protecting the butterfly.

rate through one million steps for the 16 parallel agents. However, the "PPO Only" model attained the highest reward with an 11.4% increase compared to the "GAIL + PPO" model. This may imply that the human demonstrator was imperfect in gameplay, and or the motion dynamics recorded through the Vive Tracker require a higher precision. The human demonstrator in Figure 2 attained a mean score of 48 between all three movements, which may suggest that the GAIL + PPO model successfully imitated the user to the best of their ability. While the imitation learning model did receive less reward, the GAIL + PPO model may be useful in understanding user movement bias and weakness. Personalizing agents from user demonstrations may open up pathways to autonomously adjust exercise difficulty around user day-to-day movement capabilities. Subsequently, a future evaluation must be done with a more significant amount of users to understand the ability for personalization and tuning user movement with GAIL as a reward parameter for training.

For the PPO Only model, the deep reinforcement learning alone demonstrated that PPO is highly capable of learning exercise movements by protecting the butterfly. When comparing the results of Figure 2 to the Reacher Agent reported by Juliani et al on the Unity ML_Agents Toolkit, the PPO Only model for Project IB received a 41.2% increase in cumulative reward [36]. This may suggest that games like PBF may be an ideal environment for utilizing double-jointed movements, as it was designed for upper-extremity exercise by Elor et al [28]. With the training done, the double-jointed arm for Project IB was then used to provide visual guidance for iVR exercise with PBF. Guidance was done by overlaying the IB Agent as a transparent "ghost arm" as shown in Figure 5. With the agents successfully trained, we moved on to perform a small pilot study to see how the PPO Only model competed with human agents.

## III. User Study

For this study's scope, we sought to explore how our trained PPO agent would compare to human players. Four users from the University of California Santa Cruz were recruited to compete against the trained "PPO only" model in PBF. Participants were adult college students from UCSC (one female, three males, with a mean age of 23.5 years old and 1.73 age standard deviation). Each exercise was played for one minute at ten repetitions per minute. A score point is awarded for every crystal the user blocks with the bubble shield on the butterfly. A research administrator was always present to monitor user experience and followed a strict written protocol when interacting with users. Specifically, user testing sessions consisted of the following protocol steps:

1) Preparation: The study administrator sanitized the iVR equipment, made sure all equipment was fully charged, and personally ran a session of Project IB to check the quality of motion capture data communication.
2) Introduction: The administrator instructed the user to remain still and relax. The user was verbally informed about the three exercise movements and the goal of protecting the butterfly. The user was then given a one minute tutorial for each exercise to protect the butterfly with the cooperative IB Agent "ghost arm." An example of this stage can be seen in Figure 5.
3) Rest: The user was instructed to relax for 90 seconds before performing the exercise with Project IB. This was done before every new exercise was administered.
4) Exercise: Users completed 60 seconds of gameplay while competing against the Project IB agent, and the user's final game score was recorded. Upon completion of one set, the Rest stage was repeated. An example of this stage can be seen in Figure 6. This stage was repeated until the user successfully completed all three exercises during competition with the agent.

## IV. Results and Discussion

Each of the four users from the pilot user study successfully competed with the Project IB agent. The resulting final scores between the users and agent can be seen in Table I. The Project IB agent was able to complete exercises just as well (and even slightly better) than the users for the Horizontal Shoulder Rotation movements. Nevertheless, gameplay indicated that the users were able slightly to outperform the agent for the Forward Arm Raise and Side Arm Raise exercises. Side arm raise appeared to have the highest standard deviation for the agent and the users, indicating a mixed performance. All users reported that they felt the movements were "tiring" at the speed of ten repetitions per minute (requiring a slow and controlled movement in following the butterfly).

While the initial results of Project IB were promising, there are many limitations to consider. More users must compete with both the "PPO Only" and the "PPO + GAIL" models to understand the efficacy of these models as well as exploring unlearned exercises. More demonstrations and imitation learning tuning parameters should be explored with GAIL, such that

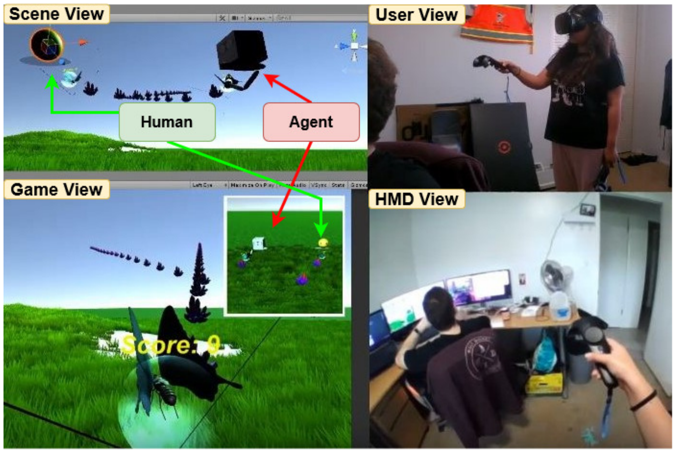


Fig. 6. Project IB Competitive Gameplay with Trained Agent. The user competes with the Project IB agent to collect the most crystals while protecting the butterfly. The agent is set to the right of the user and is tasked with protecting it's own butterfly. Crystal paths and human vs agent avatar representation are shown in the scene and game view.

| Exercise | User Score | Agent Score |
|---|---|---|
| Horizontal Shoulder Rotation | 46.6 (1.15) | **47.3 (0.58)** |
| Forward Arm Raise | **45.6 (0.58)** | 44.0 (1.00) |
| Side Arm Raise | **33.3 (4.04)** | 31.0 (1.73) |

TABLE I

RESULTS IN [MEAN (STANDARD DEVIATION)] FORMAT FOR HUMAN VERSUS AGENT GAMEPLAY. USERS WERE ADULT COLLEGE STUDENTS FROM UCSC (N=4, F=1, M=3, AGE=23.5 +/- 1.73). EACH EXERCISE WAS PLAYED FOR ONE MINUTE AT 10 REPS PER MINUTE. ONE SCORE POINT IS AWARDED PER EVERY CRYSTAL THE USER BLOCKS WITH THE BUBBLE SHIELD ON THE BUTTERFLY.

each model is tailored to each user's movement capabilities for a normalized comparison. Furthermore, a more in-depth investigation must be done to understand the effects of the cooperative "ghost arm" agent to examine if it is assistive from a presence, immersion, embodiment, and self-reported performance perspective. For example, how does the ghost arm compare to the visual guidance from crystals or no guidance at all? These limitations are being considered for future studies with our pilot data in mind.

## V. CONCLUSION

Through this paper, we presented a novel game mechanic for iVR exercise games that employed deep reinforcement learning and immersive virtual environments to learn from and help guide double-jointed exercise movements. We demonstrated how to convert a previously explored iVR exercise game for machine learning agents. We showcased a methodology of utilizing Generative Adversarial Imitation Learning and Proximal Policy Optimization to exercise with virtual butterflies. We examined two differing models for training our agents, with and without imitation learning. We demonstrated a promising learning rate through training 16 agents in parallel throughout one million steps. We evaluated one of the trained models with a set of four young adults to explore competitive applications with the agent as a game mechanic. The results suggest that with the right training parameters, the model can compete with and adhere to human-level performance in iVR for some exercises after a single training session.

In the future, we hope to explore unlearned exercises and validate a greater range of deep learning models through more extensive user testing to examine its effects on user performance, immersion, and self-reported perception. Our long term goal is to develop an at-home recovery game that uses machine learning to adapt exercise difficulty and assistance. Subsequently, we plan to explore more machine learning algorithms and input parameters such as biofeedback and musculoskeletal simulation to inform of gameplay progression. The incorporation of predictive runtime models to identify muscle weaknesses may further aid in custom movements for an individual user to help maximize their exercise by ensuring the targeted muscles are being used for a given movement. To this end, there are more butterflies to learn from as we continue working towards achieving greater physical intelligence.

## ACKNOWLEDGMENT

We thank Professor Angus Forbes of UC Santa Cruz for his advice during this project and the many participants who volunteered for this study.

## REFERENCES

[1] L. M. Howden and J. A. Meyer, *Age and sex composition, 2010*. US Department of Commerce, Economics and Statistics Administration, US . . . , 2011.

[2] CDC, "Brfss survey data and documentation 2017," C. for Disease Control, Prevention *et al.*, Eds., 2017.

[3] H. Sandler, *Inactivity: physiological effects*. Elsevier, 2012.

[4] P. Z. Pearce, "Exercise is medicine™," *Current sports medicine reports*, vol. 7, no. 3, pp. 171–175, 2008.

[5] D. Corbetta, F. Imeri, and R. Gatti, "Rehabilitation that incorporates virtual reality is more effective than standard rehabilitation for improving walking speed, balance and mobility after stroke: a systematic review," *Journal of physiotherapy*, vol. 61, no. 3, pp. 117–124, 2015.

[6] H. Mousavi Hondori and M. Khademi, "A review on technical and clinical impact of microsoft kinect on physical therapy and rehabilitation," *Journal of Medical Engineering*, vol. 2014, 2014.

[7] A. Elor, M. Teodorescu, and S. Kurniawan, "Project star catcher: A novel immersive virtual reality experience for upper limb rehabilitation," *ACM Transactions on Accessible Computing (TACCESS)*, vol. 11, no. 4, p. 20, 2018.

[8] H. G. Hoffman, W. J. Meyer III, M. Ramirez, L. Roberts, E. J. Seibel, B. Atzori, S. R. Sharar, and D. R. Patterson, "Feasibility of articulated arm mounted oculus rift virtual reality goggles for adjunctive pain control during occupational therapy in pediatric burn patients," *Cyberpsychology, Behavior, and Social Networking*, vol. 17, no. 6, pp. 397–401, 2014.

[9] H. G. Hoffman, G. T. Chambers, W. J. Meyer, L. L. Arceneaux, W. J. Russell, E. J. Seibel, T. L. Richards, S. R. Sharar, and D. R. Patterson, "Virtual reality as an adjunctive non-pharmacologic analgesic for acute burn pain during medical procedures," *Annals of Behavioral Medicine*, vol. 41, no. 2, pp. 183–191, 2011.

[10] P. J. Standen and D. J. Brown, "Virtual reality in the rehabilitation of people with intellectual disabilities," *Cyberpsychology & behavior*, vol. 8, no. 3, pp. 272–282, 2005.

[11] J. Diemer, G. W. Alpers, H. M. Peperkorn, Y. Shiban, and A. Mühlberger, "The impact of perception and presence on emotional reactions: a review of research in virtual reality," *Frontiers in psychology*, vol. 6, 2015.

[12] J. Crosbie, S. Lennon, J. Basford, and S. McDonough, "Virtual reality in stroke rehabilitation: still more virtual than real," *Disability and rehabilitation*, vol. 29, no. 14, pp. 1139–1146, 2007.

[13] P. J. Costello, *Health and safety issues associated with virtual reality: a review of current literature*. Advisory Group on Computer Graphics, 1997.

[14] M. Beccue and C. Wheelock, "Research report: Virtual reality for consumer markets," Tractica Research, Tech. Rep., Q4 2016. [Online]. Available: https://www.tractica.com/research/virtual-reality-for-consumer-markets/

[15] G. N. Yannakakis and J. Togelius, "A panorama of artificial and computational intelligence in games," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 7, no. 4, pp. 317–335, 2014.

[16] J. Fürnkranz, "Machine learning in games: A survey," *Machines that learn to play games*, pp. 11–59, 2001.

[17] T. Conde, W. Tambellini, and D. Thalmann, "Behavioral animation of autonomous virtual agents helped by reinforcement learning," in *International Workshop on Intelligent Virtual Agents*. Springer, 2003, pp. 175–180.

[18] D.-W. Huang, G. Katz, J. Langsfeld, R. Gentili, and J. Reggia, "A virtual demonstrator environment for robot imitation learning," in *2015 IEEE International Conference on Technologies for Practical Robot Applications (TePRA)*. IEEE, 2015, pp. 1–6.

[19] S.-C. Yeh, M.-C. Huang, P.-C. Wang, T.-Y. Fang, M.-C. Su, P.-Y. Tsai, and A. Rizzo, "Machine learning-based assessment tool for imbalance and vestibular dysfunction with virtual reality rehabilitation system," *Computer methods and programs in biomedicine*, vol. 116, no. 3, pp. 311–318, 2014.

[20] A. Borrego, J. Latorre, M. Alcañiz, and R. Llorens, "Comparison of oculus rift and htc vive: feasibility for virtual reality-based exploration, navigation, exergaming, and rehabilitation," *Games for health journal*, vol. 7, no. 3, pp. 151–156, 2018.

[21] S. M. Palaniappan and B. S. Duerstock, "Developing rehabilitation practices using virtual reality exergaming," in *2018 IEEE International Symposium on Signal Processing and Information Technology (ISSPIT)*. IEEE, 2018, pp. 090–094.

[22] F. Soffel, M. Zank, and A. Kunz, "Postural stability analysis in virtual reality using the htc vive," in *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology*. ACM, 2016, pp. 351–352.

[23] D. C. Niehorster, L. Li, and M. Lappe, "The accuracy and precision of position and orientation tracking in the htc vive virtual reality system for scientific research," *i-Perception*, vol. 8, no. 3, p. 2041669517708205, 2017.

[24] H. K. Kim, J. Park, Y. Choi, and M. Choe, "Virtual reality sickness questionnaire (vrsq): Motion sickness measurement index in a virtual reality environment," *Applied ergonomics*, vol. 69, pp. 66–73, 2018.

[25] T. Zhang, Z. McCarthy, O. Jow, D. Lee, X. Chen, K. Goldberg, and P. Abbeel, "Deep imitation learning for complex manipulation tasks from virtual reality teleoperation," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1–8.

[26] I. Kastanis and M. Slater, "Reinforcement learning utilizes proxemics: An avatar learns to manipulate the position of people in immersive virtual reality," *ACM Transactions on Applied Perception (TAP)*, vol. 9, no. 1, pp. 1–15, 2012.

[27] A. Rovira and M. Slater, "Reinforcement learning as a tool to make people move to a specific location in immersive virtual reality," *International Journal of Human-Computer Studies*, vol. 98, pp. 89–94, 2017.

[28] A. Elor, S. Lessard, M. Teodorescu, and S. Kurniawan, "Project butterfly: Synergizing immersive virtual reality with actuated soft exosuit for upper-extremity rehabilitation," in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 2019, pp. 1448–1456.

[29] A. Elor, S. Kurniawan, and M. Teodorescu, "Towards an immersive virtual reality game for smarter post-stroke rehabilitation," in *2018 IEEE International Conference on Smart Computing (SMARTCOMP)*. IEEE, 2018, pp. 219–225.

[30] A. Elor, M. Powell, E. Mahmoodi, N. Hawthorne, M. Teodorescu, and S. Kurniawan, "On shooting stars: Comparing cave and hmd immersive virtual reality exergaming for adults with mixed ability," *ACM Transactions on Computing for Healthcare*.

[31] A. Elor and A. Song, "isam: Personalizing an artificial intelligence model for emotion with pleasure-arousal-dominance in immersive virtual reality," in *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)(FG)*, pp. 583–587.

[32] Unity Technologies, "Unity real-time development platform — 3d, 2d vr ar," *Internet: https://unity.com/ [Jun. 06, 2019]*, 2019.

[33] HTC-Corporation, "Vive vr system," *Vive*, November 2018, https://www.vive.com/us/product/vive-virtual-reality-system/.

[34] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.

[35] M. Lanham, *Learn Unity ML-Agents–Fundamentals of Unity Machine Learning: Incorporate new powerful ML algorithms such as Deep Reinforcement Learning for games*. Packt Publishing Ltd, 2018.

[36] A. Juliani, V.-P. Berges, E. Vckay, Y. Gao, H. Henry, M. Mattar, and D. Lange, "Unity: A general platform for intelligent agents," *arXiv preprint arXiv:1809.02627*, 2018.

[37] J. M. Burnfield, K. R. Josephson, C. M. Powers, and L. Z. Rubenstein, "The influence of lower extremity joint torque on gait characteristics in elderly men," *Archives of physical medicine and rehabilitation*, vol. 81, no. 9, pp. 1153–1157, 2000.

[38] L. Ballaz, M. Raison, C. Detrembleur, G. Gaudet, and M. Lemay, "Joint torque variability and repeatability during cyclic flexion-extension of the elbow," *BMC sports science, medicine and rehabilitation*, vol. 8, no. 1, p. 8, 2016.

[39] A. K. Gillawat and H. J. Nagarsheth, "Human upper limb joint torque minimization using genetic algorithm," in *Recent Advances in Mechanical Engineering*. Springer, 2020, pp. 57–70.

[40] K. Kiguchi and Y. Hayashi, "An emg-based control for an upper-limb power-assist exoskeleton robot," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 42, no. 4, pp. 1064–1071, 2012.

[41] D. H. Perrin, R. J. Robertson, and R. L. Ray, "Bilateral isokinetic peak torque, torque acceleration energy, power, and work relationships in athletes and nonathletes," *Journal of Orthopaedic & Sports Physical Therapy*, vol. 9, no. 5, pp. 184–189, 1987.

[42] J. Hamill and K. M. Knutzen, *Biomechanical basis of human movement*. Lippincott Williams & Wilkins, 2006.

[43] M. T. Farrell and H. Herr, "Angular momentum primitives for human turning: Control implications for biped robots," in *Humanoids 2008-8th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2008, pp. 163–167.

[44] S. M. Bruijn, P. Meyns, I. Jonkers, D. Kaat, and J. Duysens, "Control of angular momentum during walking in children with cerebral palsy," *Research in developmental disabilities*, vol. 32, no. 6, pp. 2860–2866, 2011.

[45] C. Nott, R. R. Neptune, and S. Kautz, "Relationships between frontal-plane angular momentum and clinical balance measures during post-stroke hemiparetic walking," *Gait & posture*, vol. 39, no. 1, pp. 129–134, 2014.

[46] R. R. Neptune and C. P. McGowan, "Muscle contributions to whole-body sagittal plane angular momentum during walking," *Journal of biomechanics*, vol. 44, no. 1, pp. 6–12, 2011.

[47] M. Ora Powell, A. Elor, M. Teodorescu, and S. Kurniawan, "Openbutterfly: Multimodal rehabilitation analysis of immersive virtual reality for physical therapy," *American Journal of Sports Science and Medicine*, vol. 8, no. 1, pp. 23–35, 2020.

[48] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[49] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Advances in neural information processing systems*, 2016, pp. 4565–4573.