# Christian E. López

Mem. ASME
Department of Computer Science,
Afiliated to Mechanical Engineering,
Lafayette College,
Easton, PA 18042
e-mail: lopezbec@lafayette.edu

# **James Cunningham**

Mem. ASME
Department of Mechanical Engineering,
Carnegie Mellon University,
5000 Forbes Avenue,
Pittsburgh, PA 15213
e-mail: jamescun@andrew.cmu.edu

## Omar Ashour

Department of Industrial Engineering, The Pennsylvania State University, Erie, PA 16563 e-mail: oma110@psu.edu

# Conrad S. Tucker

Mem. ASME
Professor
Department of Mechanical Engineering,
Courtesy Appointment, Machine Learning,
Carnegie Mellon University,
5000 Forbes Avenue,
Pittsburgh, PA 15213
e-mail: conradt@andrew.cmu.edu

# Deep Reinforcement Learning for Procedural Content Generation of 3D Virtual Environments

This work presents a deep reinforcement learning (DRL) approach for procedural content generation (PCG) to automatically generate three-dimensional (3D) virtual environments that users can interact with. The primary objective of PCG methods is to algorithmically generate new content in order to improve user experience. Researchers have started exploring the use of machine learning (ML) methods to generate content. However, these approaches frequently implement supervised ML algorithms that require initial datasets to train their generative models. In contrast, RL algorithms do not require training data to be collected a priori since they take advantage of simulation to train their models. Considering the advantages of RL algorithms, this work presents a method that generates new 3D virtual environments by training an RL agent using a 3D simulation platform. This work extends the authors' previous work and presents the results of a case study that supports the capability of the proposed method to generate new 3D virtual environments. The ability to automatically generate new content has the potential to maintain users' engagement in a wide variety of applications such as virtual reality applications for education and training, and engineering conceptual design. [DOI: 10.1115/1.4046293]

Keywords: human computer interactions, machine learning for engineering applications, virtual and augmented reality

## 1 Introduction

The objective of procedural content generation (PCG) methods is to automatically generate content. Since the 1980s, the gaming industry has been using PCG methods to generate new game levels by manipulating game design elements such as terrains, maps, and objects [1]. Similarly, researchers have started exploring how automatically generating new content for e-learning applications can help advance Adaptive Instructional Systems (AISs), such as intelligent tutoring systems [2,3]. The ability to automatically generate new content offers several advantages for the design and development of a wide range of applications [4]. For example, automatically generating new content can help reduce the resources needed to develop new applications. PCG methods can help designers explore the design space and potentially help co-create more creative content. More importantly, content that is automatically generated can be personalized to individual's unique attributes in order to maximize the user experience [5-7]. The use of PCG methods to generate new content has been shown to improve user experience and engage users (e.g., replay value) [7–9].

In recent years, researchers have started integrating machine learning (ML) algorithms to automatically generate new content [1,10,11]. However, PCG methods that implement ML algorithms require datasets to train their generative models since these algorithms frequently use supervised learning methods. In contrast, deep reinforcement learning (DRL)-based methods are capable of generating efficient representations of complex situations and tasks by implementing sensory input information obtained from simulation environments (e.g., pixels acquired from images of a video game) [12]. Hence, there is no need to capture training data a priori, which can help reduce cost [5–7].

Given the advantages of the PCG methods and the potential of RL algorithms, this work presents a PCG method based on a Deep RL approach that generates new virtual environments. Figure 1 shows an outline of this method. A Deep RL agent is presented that generates new 3D virtual environments that are validated via a 3D simulation platform. In this work, the term "virtual" represents a 3D computer-generated (virtual) environment that users can interact with. The RL agent generates new virtual environments according to individuals' preferences for the location of a subset of virtual objects. Once a new 3D virtual environment is generated, the user can interact with it using a variety of interfaces (e.g., immersive virtual reality (VR) headset, smartphone, and computer). This work extends the authors' previous work [13] and presents the results of a case study that supports the ability of the proposed method to generate new 3D virtual environments.

## 2 Literature Review

**2.1 Procedural Content Generation.** Procedural content generation can be defined as the field that studies the development of algorithms and methods capable of automatically generating content. The gaming industry has used PCG for decades [1]. Most of the early PCG methods were composed of rule sets and heuristics that guided the content generation process or functions to evaluate the generated content. These heuristics and functions were developed by designers based on their understanding of the application [6,14]. However, in recent years, researchers have started exploring the use of supervised ML algorithms to train generative models capable of automatically creating new content [11].

One of the most well-known projects that integrate supervised ML to generate new game environments is "Mario AI" [15–17].<sup>2</sup>

<sup>&</sup>lt;sup>1</sup>Corresponding author.

Manuscript received September 30, 2019; final manuscript received January 30, 2020; published online June 3, 2020. Assoc. Editor: Caterina Rizzi.

<sup>&</sup>lt;sup>2</sup>http://marioai.org/

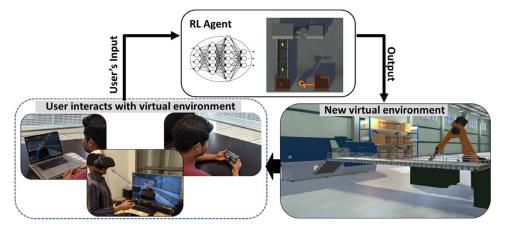


Fig. 1 Outline of the reinforcement learning PCG method

Researchers have presented a wide range of PCG methods to automatically generate new environments for a variety of popular games [1,14-17]. For example, Summerville and Mateas [1] introduce a Long Short-Term Memory Recurrent Neural Network framework to generate new Super Mario Brothers levels. Their model was trained using a corpus of 39 existing levels of Super Mario Brothers, which they were able to augment by using several training techniques (e.g., stacking). Moreover, Summerville et al. [17] present a Bayesian network to automatically generate level topologies for Zelda-like games. They annotated the topology characteristics of 38 levels of different Zelda games in order to train their generative model. Justesen et al. [18] attempt to overcome the overfitting problem that arises when training DRL agents on static game environments (e.g., training a model to play an Atari game by using just one level) by introducing a search-based PCG. Their progressive PCG method helps control for the difficulty of levels to match the Deep RL agent being trained to play the game. They trained their models using levels from the games of Zelda, Solarfox, Frogs, and Boulderdash. Similarly, researchers have introduced PCG methods based on Markov Chain [19,20], and matrix factorization approaches [21]. However, these methods still require human-authored content to train their models.

Most of the current PCG methods that implement supervised ML methods require some initial dataset to train their generative models. In contrast, RL algorithms implement high-dimensional sensory input to generate efficient representations of complex situations and tasks with the use of simulation [12]. Hence, there is no need to capture training data a priori. Based on these advantages, this work presented a PCG method to generate new 3D virtual environments based on a Deep RL approach.

**2.2 Adaptive Instructional Systems.** The field of AISs has greatly benefited from integrating methods to generate new content for their adaptive applications [22]. These types of systems require significantly more content than their non-adaptive counterparts since for each adaptation, new content is required [23]. AISs are defined as "class of intelligent, machine-based tools that guide learning experiences by tailoring instruction and recommendations based on the goals, needs, and preferences of each learner [or team] in the context of domain learning objectives" [23]. Intelligent tutoring systems, intelligent method, recommender systems, personal assistants, and intelligent instructional media fall under the umbrella of AISs.

Within this field, RL has been used to model students' learning styles and develop pedagogical policy strategies [3,9,24,25]. However, there has been a limited number of studies that have explored how to automatically generate new content for learning purposes [26,27]. For example, Hullett and Mateas [8] present an application capable of generating new scenarios for a firefighting training

application. The application was able to generate different scenarios of buildings partly collapsed based on the desired skills the users wanted to train on. Smith et al. [28] implement a method for creating levels in a learning application aimed at teaching students about fractional arithmetic. The method implements a constraint-focused generator design approach. Similarly, a learning application that implemented PCG and gamification to engage students in solving math problems is introduced in Ref. [29]. This method was founded on template-based and constructive algorithms.

In the context of conflict resolution, a serious game application that combined a player modeling and a metaheuristic-search PCG approach is introduced in Ref. [30]. This PCG method was driven by a neural network used to predict the distribution fairness of the players. The results of this study support the value of PCG to guide the learning of individuals toward targeted learning objectives. Most recently, Hooshyar et al. [7] proposed a PCG framework for educational game applications based on a genetic algorithm (GA) approach. The framework allows designers to control the generation process, given various learning objectives and preferences. In a different study, Hooshyar et al. [26] present a data-driven PCG approach based on genetic and support vector machine algorithms. They implemented their method in a language learning application and compared the method against a heuristic-based approach. Their results indicate that their data-driven approach was more effective at generating content that matched the performance target of individuals, compared to the heuristic approach. Similarly, Sottilare [23] presents an ML method based on a GA approach to automatically generate new scenarios from a set of parent scenarios for virtual instructional and game-based applications.

The previous studies show how PCG methods can be implemented in learning applications and their potential benefits. These studies also show that researchers are starting to use ML approaches (e.g., neural network, support vector machines, and genetic algorithms) to train their PCG models. They train their models on datasets from existing content or datasets containing users' data, which has to be generated or collected a priori [7,26,30]. The process of generating new content to use as a training dataset can require significant time and resources [5-7]. In recent years, researchers have started exploring how realistic, synthetic data can be automatically generated [31,32]. However, while studies have shown that these approaches can generate synthetic datasets that cannot be accurately distinguished from human-generated ones [33–35], they still require some initial datasets to train their models. In contrast, RL approaches do not require training data to be collected a priori since they take advantage of simulation to train their models. Based on the limitations of supervised ML algorithms and the advantages of RL algorithms, this work presents a PCG method based on a Deep RL approach. The RL agent is trained using a simulation platform to automatically generate new 3D virtual environments, which could potentially be used for learning applications.

**2.3 Reinforcement Learning.** While traditional supervised ML algorithms require the use of a training dataset, RL methods do not require a training dataset to be collected a priori since they take advantage of simulation environments to generate efficient representations of complex situations and tasks [12]. The RL process can be understood as a Markov decision process, where the RL agent connects to a simulation environment via different sensory inputs. The objective of the agent is to develop a model that selects the actions that maximize its long-run reward. In other words, the agent creates the desired action policy by the process of trial and error via simulation [36]. Hence, an RL agent can be described as a software agent capable of inducing an action policy in an uncertain environment with delayed rewards [37].

Reinforcement learning methods are suitable for solving learning control problems, which are challenging for traditional supervised ML algorithms and dynamic programming optimization methods [38]. RL agents focus on generating an action policy that can adapt to changes in the environment (e.g., state space). Researchers have used RL methods to train agents capable of mastering complex tasks at human-level performance [39–41]. In recent years, Deep RL algorithms have been implemented to master and perform a wide range of tasks, from Atari games to the Chinese game of Go [39,40]. Thanks to these advancements, researchers argue that these algorithms will revolutionize the field of artificial intelligence [12]. In addition, the rapid development of these RL methods has been encouraged by the rise of easy to use, scalable simulation platforms [42–44].

In the context of AISs, RL-based methods have shown promising results in helping personalized narrative-centered learning environments. For example, Wang et al. [45] present a Deep RL framework to personalize interactive narrative for an educational game. Similarly, Rowe et al. [46] introduce a multi-armed bandit computational formalism, consisting of several components of a Deep RL framework to generate a new training scenario for the Army. The authors also explored Long Short-Term Memory Network approaches and stated that in future work, they would be implementing RL algorithms to help generate new complex training scenarios.

Table 1 shows a summary of existing work related to methods that automatically generate content (i.e., PCG). This table shows that while PCG methods are frequently used in gaming applications, researchers are starting to explore the use of PCG methods for learning purposes. However, most of the studies on learning applications implement metaheuristics. In light of the advantages of PCG methods and the potential of RL algorithms, this work presents a PCG method based on an RL approach that generates new 3D virtual environments. The RL agent validates the new 3D virtual environments via a simulation platform; hence, it does not require any training data to be collected a priori. The RL-based PCG method is implemented in a case study to generate new layouts for a virtual 3D manufacturing environment used for an e-learning application.

In the authors' previous work, initial results of the performance of the RL agent's reward score were presented [13]. The results show that the RL agent did not reach the maximum reward score, but that its reward score was significantly and strongly correlated with the training iterations ( $\rho$ =0.98, p-value<0.001). In other words, the RL agent was not able to generate a 3D virtual environment that was completely functional and that maximized the rewards function. However, it managed to model an action policy that

maximized the long-run rewards function. Moreover, in the previous work, the training of the agent was not parallelized and the training time constrained to less than 6 h due to computational limitations. These factors played a significant role in the performance of the RL agent. Based on these limitations, in this work, the authors extended their previous study by implementing parallelized training over 60,000 iterations. In addition, the reward function of the RL agent and the simulation environment used for training have been enhanced in order to incentivize the generation of more realistic and functional layouts. Finally, the results of a case study that supports the capability of the proposed method to generate new 3D virtual environments are presented in this work.

## 2.4 Reinforcement Learning and Operations Research.

The objective of PCG methods to generate new environments given certain criteria is analogous to the operations research (OR) problem of facility layout planning (FLP). The objective of FLP algorithms is to identify the optimal arrangement of equipment or facilities in accordance with some criteria and given certain constraints [47]. FLP problems are an NP-complete problem, which means that "the computational time required to find an optimal solution increases exponentially with the problem size" [48]. This is one of the reasons why researchers have proposed multiple metaheuristics algorithms to solve the FLP problem, such as simulated annealing and genetic algorithms [47]. However, one of the limitations of optimization approaches is that a given optimal solution might not continue to be optimal under a different problem configuration. For example, if an additional constraint is added (e.g., now machine Z must be in the coordinates x and y), the algorithm needs to be run again to find the optimal or near-optimal solution [47,49]. In contrast, since RL algorithms focus on generating an action policy that can adapt to changes in the state space, they do not require additional training when exposed to a new state (e.g., now machine Z must be in the coordinates x and y).

Due to the advantages of RL algorithms, researchers have explored how to implement RL in combination with metaheuristics with the objective of identifying more efficient methods for solving OR problems [50–55]. Recently, some studies have shown promising results of using RL for solving combinatorial optimization problems [56]. For example, RL algorithms have been implemented to tackle classical OR problems like dynamic job shop scheduling problem [57], vehicle touting problem [58], among others routing and scheduling problems [56]. In a recent study, Govindaiah and Petty [59,60] present the application of a framework that integrates RL algorithms and discrete event simulation to improve the cost efficiency of material handling plans under varying product demands. Their method focused on reducing the cost of material handling plans by changing the routes, timing, and equipment used to transport the material between workstations and/or warehouses. However, their method did not consider the locations of the workstations nor warehouses, the reason why it cannot be implemented for FLP [49]. The case study used in this work to test the proposed Deep RL PCG method shares some characteristics with the material handling problem tackled by Govindaiah and Petty [59,60]. However, the proposed method focuses on generating new 3D virtual environments by allocating a set of virtual objects. In the case study presented, the RL agent is capable of changing the location of the workstation (i.e., injection molding machine, see Sec. 4) and material handling equipment (e.g., conveyor belts and

Table 1 Summary of the related work

Reference	Metaheuristics	Supervised ML	RL	Environment generation	Application context
[7,8,28,29]	X				Learning
[9,26,30,45]		X			Learning
[1,14–18]		X		X	Games
This work			X	X	Learning/games

robot arms). Moreover, the proposed Deep RL PCG method can adapt to changes in the problem space (i.e., state space) without the need for additional training. That is, once the RL agent is trained, it can generate new 3D virtual environments given different injection molding machine locations (see Sec. 5 for results). In contrast, using traditional OR methods would require to run optimization or metaheuristic algorithms every time the problem space changes when a constraint is added or modified (e.g., now machine Z must be in the coordinates x and y) [47,49].

#### 3 Method

In this work, a PCG method based on a Deep RL approach is introduced. The method is capable of dynamically generating new 3D virtual environments by implementing an RL agent that validates the content via a 3D simulation platform. Figure 2 shows the method of the Deep RL algorithm implemented. In addition, it shows 2D aerial views of the 3D simulation platform used to validate the virtual manufacturing environments generated for the case study (see Sec. 4).

Reinforcement learning problems are framed as Markov decision processes, where the agents connect to the simulation environment at a given time t via the sensory inputs of state  $(S_t)$  that belongs to the set of possible states S, and action  $(A_t)$  that belongs to the set of possible actions A (see Fig. 2). In each training epoch t, the agent observes the current state:  $S_t$  and chooses an action to be executed:  $A_r$ . The environment reacts to the action executed and determines the new state to transition:  $S_{t+1}$ , as well as the reward signal (i.e., reinforcement signal):  $R_t$ . The sensory inputs of the state and action can be in a vector form, containing information about the state of the environment and information regarding the action the agent is taking, respectively. The agent makes decisions based on a policy that is defined by a mapping from the state space to a probability distribution over the action space, formalized as  $\pi(S_t) \in P(A)$ . In Deep RL, this policy function is realized using a neural network which takes  $S_t$  as input and generates probabilities for selecting each possible action as output.

The goal of an RL agent is to determine a particular policy  $\pi^*$  which maximizes the long-run reward of the agent. The long-run reward, also known as the return, is used as an objective function over the reward signal itself because it is more stable and less sparse. The return is defined as  $\rho = \sum_{t=0}^T \gamma^t R_t$ , where  $\gamma \in [0,1]$  is

the discount factor that controls the exponential devaluation of delayed rewards.

In this work, the proximal policy optimization (PPO) [61] algorithm is employed to train the RL agent. PPO is a policy gradient approach to RL based on the Trust Region Policy Optimization algorithm introduced by Schulman et al.'s work [62]. Schulman et al.'s [61] study reveals that the PPO algorithm outperformed other policy gradient algorithms, and provided a more favorable tradeoff between sample complexity, simplicity, and wall time. Given that a neural network is fully differentiable, gradient ascent can be applied to the policy function directly with respect the advantage estimate of the policy, a quantity which is related to the expected value of the policy's return (readers are referred to Refs. [61,62] for additional details).

For generating a new 3D virtual environment,  $S_t$  contains enough information to describe the current state of the virtual environment (e.g., location, orientation, and relevant properties of objects in the environment), while  $A_t$  corresponds to ways that the agent can alter the environment. These actions could correspond to determining the location, orientation, or parameters which govern the behavior of the virtual objects in the 3D virtual environment (see Fig. 2). Given this framing, the proposed DRL method can be applied to generate 3D virtual environments for problems that can be framed as arranging multiple individual objects with inherent properties and that can be expressed in a vector form. This enables the use of this method in a variety of application, such as educational and training where new environments are generated for learning purposes or game applications where new levels are generated for entertainment purposes.

The goal of the RL agent is to develop a model that selects the actions that maximize its long-run reward signal, which takes the form of a scalar value. The elements of the reward function will depend on the behavior that the designers expected the RL agent to model (i.e., learn). The RL agent needs to be rewarded for generating new environments that are functional and not just a random placement of virtual objects. This can be achieved by designing a reward function that incentivizes the generation of functional environments and penalizes nonfunctional ones (e.g., makes parts as in the manufacturing layout example of Figs. 1 and 2). However, a major difference between layout generation problems and other RL problems is that each action (i.e., placement of an object) cannot be evaluated until the full layout has been generated. This means that every action except the final will have an immediate

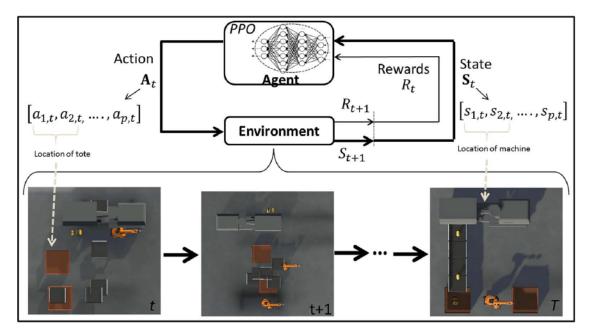


Fig. 2 Reinforcement learning framework representation

reward of  $R_t = 0$ . Thus, the proposed method omits the discount factor from the return, by choosing  $\gamma = 1$ .

Through its interactions with the simulated environment, the RL agent is trained (i.e., learns) to model an action policy that will maximize its return. Once the RL agent is trained, it will be able to generate new 3D virtual environments, given an initial state provided by the user or randomly selected by the agent. In the example shown in Fig. 2, this could be the initial location of the injection molding machine. Hence, the RL agent will place the objects in a way that will create a functional virtual 3D manufacturing layout.

#### 4 Case Study

For this case study, the authors used a VR learning application designed to teach industrial engineering (IE) concepts (i.e., Poisson distribution, Little's law, and queuing theory) with the use of a simulated manufacturing system. Specifically, a manufacturing system that produces power drills was simulated, as shown in Figs. 1 and 2. The objective of this VR learning application is to provide a tool with a common theme that educators could use to teach IE concepts and integrate course knowledge into their curriculum [63]. A power drill manufacturing line was selected since previous studies that aim to integrate IE course knowledge have implemented similar power tools [64]. The virtual environment simulates the initial steps of the process to manufacture a power drill, where the plastic housing is manufactured.

Figure 3 shows, from a user's point of view, a functional layout for this manufacturing system. In this layout, first, an injection molding press produces the plastic housing components. Then, they are cooled down with the use of a conveyor belt. Finally, the plastic housings are placed in a tote with the use of a robotic arm in order to be transported to the assembly line. The 3D virtual environment allows users to interact with virtual objects. For this application, the agent is rewarded based on the efficiency and functionality of the layout generated to produce goods (e.g., the rightmost image on Fig. 2 has a high reward score, while the two other images have a low reward score). Specifically, the reward function used in this case study can be mathematically expressed as follows:

$$R = \beta_1(\text{Tote}) - \beta_2(\text{Floor}) + \beta_3 \sigma_{\lambda} + \beta_4(\text{Flow})$$
 (1)

For

$$Tote = \sum_{p=1}^{P} \phi_p$$
 (2)

$$Floor = \sum_{p=1}^{p} \varphi_p$$
 (3)

$$\sigma_{\lambda} = \sqrt{\frac{\sum_{e=1}^{E} (\lambda_e - \lambda_U)^2}{E}}$$
 (4)

$$Flow = \sum_{p=1}^{P} \sum_{e=1}^{E} \Delta_{p,e}$$
 (5)

$$\beta_i > 0 \ \forall i = \{1, 2, 3, 4\}$$
 (6)

where

- φ<sub>p</sub> is a binary variable that indicates if a given part p was correctly placed in a tote φ<sub>p</sub> = 1 or not φ<sub>p</sub> = 0, for p ε {P}
- φ<sub>p</sub> is a binary variable that indicates if a given part p falls to the floor φ<sub>p</sub> = 1 or not φ<sub>p</sub> = 0, for p ε {P}
- λ<sub>e</sub> is a parameter that describes the behavior distribution of equipment e, for e ε {E}. (In the case study, this parameter is only applied to the conveyor and injection machine)
- Δ<sub>p,e</sub> is a binary variable that indicates if a given part p interacted with a given equipment e, for p ε {P} and e ε {E}

The reward function shown in Eq. (1) will be maximized when all the parts p are placed in a tote and no parts fall on the floor, following Eqs. (2) and (3), when the standard deviation of the parameters that describe the behavior distribution of the equipment set {E} is minimized, following Eq. (4), and when all the parts interact with all the equipment following the manufacturing process, as shown in Eq. (5). This reward function was designed to reinforce the generation of functional manufacturing layouts that follow the predefined manufacturing process and have a constant flow of part being placed in the tote. This reward function will be computed for every simulation epochs  $t(R_t)$ , as shown in Fig. 2. In addition, in every simulation epoch t, the RL agent will be able to control the placement  $(x_e, y_e)$  and the parameters that describe the behavior distribution  $(\lambda_e)$  of the equipment set  $\{E\}$ . The environment will provide the agent with the state information about the placement  $(x_u, y_u)$  of the equipment placed by the user  $\{U\}$ . The set of equipment {U} will allow users to customize the VR environment. In the event that the user does not need to customize the environment, the equipment set {U} can be placed randomly to generate a new environment.

For this application, users can select the location of the injection molding machine,  $U = \{\text{injection molding machine}\}\$ . On the other hand, the RL agent will manipulate one conveyor belt, one tote, and one robot arm. This means that the set E will contain three different equipment (i.e., virtual objects). In order for the agent to be robust to various placements of the injection molding machine, the position of the injection machine was randomly changed at every epoch t, and thus to maximize the reward, the agent would need to generate a functional layout regardless of the position of the machine. The  $\lambda_{conveyor}$  parameter will control the speed of the conveyor, while the  $\lambda_{\text{machine}}$  parameter will control the speed of the injection molding machine. To improve training performance, the RL agent is trained in parallel with multiple layouts. This provides the benefit in improved diversity of training samples by ensuring that the agent is exploring multiple action trajectories simultaneously. In this case study, the agent is trained in parallel on 32 environments.



Fig. 3 User's point of view of a functional manufacturing layout

Finally, in this work, the game engine Unity [65] is used as the 3D simulation platform to train the RL agent. Because of its fidelity, physics simulation capabilities, accessibility, and community support, Unity is widely used by developers [66,67], as well as by researchers [33,68]. Furthermore, Unity ML-Agents Toolkit [44] provides several algorithms and functionalities for the development and design of RL-based applications [69]. In this case study, for each simulation epoch t, a total of 10 parts were simulated (i.e.,  $p = \{1-10\}$ ). This number of parts was selected to reduce the complexity of the simulation while allowing the simulation platform to generate the state transition in which the environment reacts to the action executed and provides a reward signal. However, this number can be increased, and the relative difference between the rewards score of layouts would not change. That is, a layout that allows all the parts to fall on the floor will always have a worse reward score than one that places all the parts on the tote, no matter how many parts are simulated.

#### 5 Results and Discussions

The RL agent was trained using an Intel®  $Core^{TM}$  i7-4770 K 3.50 GHz CPU and 16 GB RAM. A total of 10,000 training iterations (t=10,000) on 32 simulated environments were used to train an RL agent in parallel. This means that a total of 32,000 virtual environments were generated and evaluated to train the RL agent. The total training time was 3.25 h. In this work, the coefficient of our rewards function (i.e.,  $\beta_1$ ,  $\beta_2$ ,  $\beta_3$ ,  $\beta_4$ ) were empirically set to one in order to give the same importance to the elements of the reward function. Figure 4 shows the evolution of the RL agent's average reward over the 32 environments, given the training epoch t. The y-axis shows the bounds of the reward function (i.e., [-11, 21]). Figure 4 shows that the agents' rewards score was significantly and strongly correlated with the

simulation iterations ( $\rho$  = 0.915, p-value < 0.001). This indicates that the agent managed to train a model that describes an action policy that maximized the long-run rewards function used in this case study.

To test the performance of the trained RL agent, a total of 512 (i.e., 32\*16) new 3D virtual environments were generated and evaluated. This process took 2 min and 20 s. That is, the trained RL agent takes, on average, 0.27 s to generate a new 3D virtual environment. Table 2 shows the number of layouts generated, given the rewards score achieved and a description as to why the rewards score was not optimal (i.e., 21). This table shows that, on average, these layouts had a reward score of 13.18 (SD=7.74). This is in line with the average rewards score achieved during the last iteration of the training process (see Fig. 4). It also shows that more than 55.08% of the layouts achieved a reward score greater than 18 and only mismatched the speed between the conveyor and the injection molding machine (i.e.,  $\sigma_{\lambda}$ ). Figure 5 shows 2D aerial views of several of the environments generated for testing the performance of the trained RL agent.

The results indicate that the agent managed to train a model that describes an action policy that maximized the long-run rewards function used in this case study. Moreover, the results show that the trained RL agent is capable of generating new 3D virtual environments given different injection molding machine locations without the need for additional training and in less than a second. This is in contrast with common methods used for the FLP, which required to rerun optimization algorithms every time the problem space changes (e.g., the injection molding machine location changes). This finding shows promising results for using PCG methods based on Deep RL approaches to generate new 3D virtual environments. The capability to generate new 3D virtual environments given different initial configurations of virtual objects can help personalize applications to an individual's unique preferences.

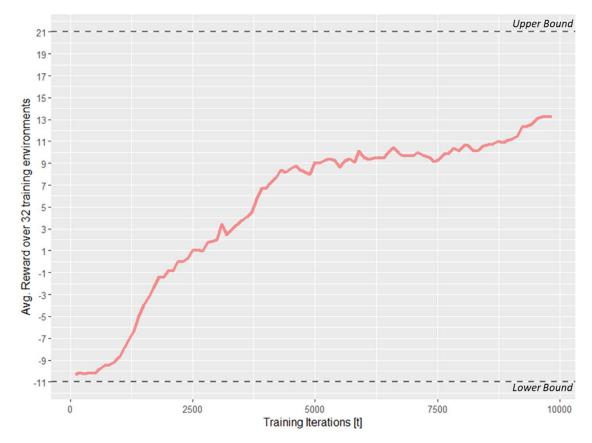


Fig. 4 Reinforcement learning agent rewards score versus training iterations

Table 2 Summary of environments generated to evaluate the trained model

Reward	No.	Percentage	Comments
21	45	8.79	Optimal layout
19	237	46.29	Mismatched speed only
11	25	4.88	No robot or conveyor interaction, but matched speed and all parts in the tote
9	128	25.00	No robot or conveyor interaction, mismatched speed, but all parts in the tote
1	11	2.15	Robot and conveyor interaction, matched speed, but no parts in the tote
-1	52	10.16	Robot and conveyor interaction but no parts in bin, mismatched speed
-4	3	0.59	Conveyor interaction and matched speed, but no robot interaction, nor parts in bin
-6	9	1.76	Conveyor interaction but no robot interaction, no parts in tote, and mismatched speed
-11	2	0.39	No conveyor or robot interaction, mismatched speed, and no parts in the tote



Fig. 5 Example of new manufacturing layouts generated by the trained RL agent. Eight randomly generated layouts from the learned policy. The reward achieved by each layout is displayed underneath. Rewards of 21 and 19 reflect the use of all objects in the layout to place all ten parts in the bin. Rewards of 11 and 9 reflect that all ten parts were placed in the bin, but the conveyor belt and robot arm were not utilized. The difference within these categories reflects whether the speed of the conveyor matches the speed of the injection machine.

# 6 Conclusion and Future Works

The ability to automatically generate new content with the use of PCG methods offers several advantages for the development and design of new applications. PCG methods can help reduce the resources needed to develop new applications. More importantly, content that is automatically generated can be personalized and adapted to an individual. Implementing PCG methods allows designers to generate new environments that can help improve the overall user experience. Researchers have started developing PCG methods that integrate supervised ML algorithms, which allow designers to generate new content more efficiently compare to heuristics-based methods. However, these algorithms require large datasets to train their generative models. In contrast, RL methods do not require any training data to be collected a priori since they take advantage of simulation environments to generate efficient representations of complex situations and tasks.

In light of this, a PCG method based on a Deep RL approach that generates new virtual environments is presented. This method trains

a model by implementing an RL agent that validates new 3D virtual environments via a 3D simulation platform; hence, it does not require any training data to be collected a priori. In this work, a case study is introduced where the proposed method is used to generate new 3D virtual manufacturing environments, with the intention to teach IE concepts. The preliminary results indicate that the RL agent was able to model (i.e., learn) a policy that allows it to automatically generate new and functional 3D virtual environments.

The proposed Deep RL PCG approach can help designers automatically generate new content for a wide range of applications. For example, Fig. 6 shows how the 3D virtual environment generated for the case study can be integrated into an immersive VR learning application. This immersive VR application can help users learn about IE concepts. The PCG method presented can also be applied to other applications that can benefit from automatically generating new 3D virtual environments (e.g., Adaptive Instructional Systems, adaptive games). Designers can implement this method in their applications by creating a rewards function based on the new environments they would like the RL agent to generate.

While this work presents a novel PCG method based on a Deep RL approach, there still exist a lot of areas for improvement. First, the method should be used to generate other types of 3D virtual environments that differ from the manufacturing layouts used in the case study. Moreover, while the RL approach does not require the collection of data a priori since it takes advantage of simulation to train its model, the reward function, which impacts the action policy the RL agent models, can be challenging to design under certain conditions. Furthermore, it could be challenging under certain circumstances to create simulation environments that allow an RL agent to identify the desired action policy. More importantly, future work should explore how integrating the proposed PCG method into learning applications can impact the motivation and learning of users. For example, as shown in Fig. 5, this method can be used to generate new 3D virtual environments for immersive VR learning applications. However, the impact of automatically generating new content on users' learning and engagement still has to be tested. Nevertheless, this work presents initial groundwork on integrating RL algorithms to automatically generate new content, which has significant implications for personalized and adaptive systems.



Fig. 6 Users interacting with the generated virtual environment using an immersive VR headset

#### Acknowledgment

This research is funded by the National Science Foundation NSF DUE #1834465. Any opinions, findings, or conclusions found in this paper are those of the authors and do not necessarily reflect the views of the sponsors. The authors would also like to thank the hard work of Bradley Nulsen, Gerard Pugliese Jr., Adith Rai, and Matthew Rodgers on developing and implementing the application used in this work.

#### References

- Summerville, A., and Mateas, M., 2016, "Super Mario as a String: Platformer Level Generation Via LSTMs," arXiv preprint, arXiv:1603.00930.
- [2] Iglesias, A., Martínez, P., and Fernández, F., 2003, "Navigating Through the RLATES Interface: A Web-based Adaptive and Intelligent Educational System," OTM Confederated International Conferences On the Move to Meaningful Internet Systems, Springer, Berlin, Heidelberg, pp. 175–184.
- [3] Fenza, G., Orciuoli, F., and Sampson, D. G., 2017, "Building Adaptive Tutoring Model Using Artificial Neural Networks and Reinforcement Learning," Proceedings—IEEE 17th International Conference on Advanced Learning Technologies (ICALT), Timisoara, Romania, IEEE, pp. 460–462.
- [4] Togelius, J., Yannakakis, G. N., Stanley, K. O., and Browne, C., 2011, "Search-Based Procedural Content Generation: A Taxonomy and Survey," IEEE Trans. Comput. Intell. AI Games, 3(3), pp. 172–186.
- [5] Bidarra, R., de Kraker, K. J., Smelik, R. M., and Tutenel, T., 2010, "Integrating Semantics and Procedural Generation: Key Enabling Factors for Declarative Modeling of Virtual Worlds," Proceedings of the FOCUS K3D Conference on Semantic 3D Media and Content, France.
- [6] Smith, G., Whitehead, J., and Mateas, M., 2011, "Tanagra: Reactive Planning and Constraint Solving for Mixed-Initiative Level Design," IEEE Trans. Comput. Intell. AI Games, 3(3), pp. 201–215.
- [7] Hooshyar, D., Yousefi, M., and Lim, H., 2018, "A Procedural Content Generation-Based Framework for Educational Games: Toward a Tailored Data-Driven Game for Developing Early English Reading Skills," J. Educ. Comput. Res., 56(2), pp. 293–310.
- [8] Hullett, K., and Mateas, M., 2009, "Scenario Generation for Emergency Rescue Training Games," Proceedings of the 4th International Conference on Foundations of Digital Games—FDG '09, Orlando, FL, pp. 99–106.
- [9] Sawyer, R., Rowe, J., and Lester, J., 2017, "Balancing Learning and Engagement in Game-Based Learning Environments with Multi-Objective Reinforcement Learning," International Conference on Artificial Intelligence in Education, Springer, Cham, pp. 323–334.
- [10] Hendrikx, M., Meijer, S., Van Der Velden, J., and Iosup, A., 2013, "Procedural Content Generation for Games: A Survey," ACM Trans. Multimed. Comput. Commun. Appl., 9(1), pp. 1–22.
- [11] Yannakakis, G. N., 2012, "Game AI Revisited," Proceedings of the 9th Conference on Computing Frontiers—CF '12, Cagliari, Italy, pp. 285–292.
  [12] Arulkumaran, K., Deisenroth, M. P., Brundage, M., and Bharath, A. A., 2017,
- [12] Arulkumaran, K., Deisenroth, M. P., Brundage, M., and Bharath, A. A., 2017, "Deep Reinforcement Learning: A Brief Survey," IEEE Signal Process. Mag., 34(6), pp. 26–38.
- [13] Lopez, C. E., Ashour, O., and Tucker, C. S., 2019, "Reinforcement Learning Content Generation for Virtual Reality Applications," International Design Engineering Technical Conference and Computers and Information in Engineering Conference (DETC2019-97711), Anaheim, CA, August.
- [14] Shaker, N., Togelius, J., Yannakakis, G. N., Weber, B., Shimizu, T., Hashiyama, T., Sorenson, N., Pasquier, P., Mawhorter, P., Takahashi, G., Smith, G., and Baumgarten, R., 2011, "The 2010 Mario AI Championship: Level Generation Track," IEEE Trans. Comput. Intell. AI Games, 3(4), pp. 332–347.
- [15] Shi, P., and Chen, K., 2018, "Learning Constructive Primitives for Real-Time Dynamic Difficulty Adjustment in Super Mario Bros," IEEE Trans. Comput. Intell. AI Games, 10(2), pp. 155–169.
- [16] Guzdial, M. J., Sturtevant, N., and Li, B., 2016, "Deep Static and Dynamic Level Analysis: A Study on Infinite Mario," Twelfth Artificial Intelligence and Interactive Digital Entertainment Conference, San Francisco, CA, pp. 31–38.
- [17] Summerville, A. J., Behrooz, M., Mateas, M., and Jhala, A., 2015, "The Learning of Zelda: Datadriven Learning of Level Topology," Proceedings of the FDG Workshop on Procedural Content Generation in Games, Pacific Grove, CA, pp. 31–36.
- [18] Justesen, N., Torrado, R. R., Bontrager, P., Khalifa, A., Togelius, J., and Risi, S., 2018, "Illuminating Generalization in Deep Reinforcement Learning Through Procedural Level Generation," arXiv preprint arXiv:1806.10729.
- [19] Dahlskog, S., Togelius, J., and Nelson, M. J., 2014, "Linear Levels Through n-Grams," Proceedings of the 18th International Academic MindTrek Conference on Media Business, Management, Content & Services, Tampere, Finland, pp. 200–206.
- [20] Snodgrass, S., and Ontañón, S., 2013, "Generating Maps Using Markov Chains," Ninth Artificial Intelligence and Interactive Digital Entertainment Conference: Papers From the 2013 AIIDE Workshop, Boston, MA, pp. 25–28.
- [21] Shaker, N., and Abou-Zleikha, M., 2014, "Alone We Can Do So Little, Together We Can Do So Much: A Combinatorial Approach for Generating Game Content," AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, Raleigh, NC, pp. 167–173.
- [22] Almohammadi, K., Hagras, H., Alghazzawi, D., and Aldabbagh, G., 2017, "A Survey of Artificial Intelligence Techniques Employed for Adaptive Educational

- Systems Within e-Learning Platforms," J. Artif. Intell. Soft Comput. Res., 7(1), pp. 47-64.
- [23] Sottilare, R. A., 2018, "A Hybrid Machine Learning Approach to Automated Scenario Generation (ASG) to Support Adaptive Instruction in Virtual Simulations and Games," 8th International Defense and Homeland Security Simulation Workshop, Budapest, Hungary, pp. 1–9.
- [24] Dorça, F. A., Lima, L. V., Fernandes, M. A., and Lopes, C. R., 2013, "Comparing Strategies for Modeling Students Learning Styles Through Reinforcement Learning in Adaptive and Intelligent Educational Systems: An Experimental Analysis," Expert Syst. Appl., 40(6), pp. 2091–2101.
- [25] Iglesias, A., Martínez, P., Aler, R., and Fernández, F., 2009, "Learning Teaching Strategies in an Adaptive and Intelligent Educational System Through Reinforcement Learning," Appl. Intell., 31(1), pp. 89–106.
- [26] Hooshyar, D., Yousefi, M., Wang, M., and Lim, H., 2018, "A Data-Driven Procedural-Content-Generation Approach for Educational Games," J. Comput. Assist. Learn., 34(6), pp. 731–739.
- [27] Hooshyar, D., Yousefi, M., and Lim, H., 2017, "A Systematic Review of Data-Driven Approaches in Player Modeling of Educational Games," Artif. Intell. Rev., 52(3), pp. 1–27.
- [28] Smith, A. M., Andersen, E., Mateas, M., and Popović, Z., 2012, "A Case Study of Expressively Constrainable Level Design Automation Tools for a Puzzle Game," Proceedings of the International Conference on the Foundations of Digital Games, New York, NY, pp. 156–163.
- [29] Rodrigues, L., Bonidia, R. P., and Brancher, J. D., 2017, "A Math Educational Computer Game Using Procedural Content Generation," Brazilian Symposium on Computers in Education (Simpósio Brasileiro de Informática na Educação-SBIE), Vol. 28, No. 1, p. 756.
- [30] Grappiolo, C., Cheong, Y. G., Togelius, J., Khaled, R., and Yannakakis, G. N., 2011, "Towards Player Adaptivity in a Serious Game for Conflict Resolution," Proceedings—2011 3rd International Conference on Games and Virtual Worlds for Serious Application, Athens, Greece, pp. 192–198.
- [31] Dering, M. L., and Tucker, C. S., 2017, "Generative Adversarial Networks for Increasing the Veracity of Big Data," Proceedings—2017 IEEE International Conference on Big Data, Boston, MA, IEEE, pp. 2595–2602.
- [32] Beecks, C., Uysal, M. S., and Seidl, T., 2015, "Gradient-Based Signatures for Big Multimedia Data," IEEE International Conference on Big Data, Santa Clara, CA, IEEE, pp. 2834–2835.
- [33] Lopez, C. E., Miller, S. R., and Tucker, C. S., 2019, "Exploring Biases Between Human and Machine Generated Designs," ASME J. Mech. Des., 141(2), p. 021104.
- [34] Lopez, C., Miller, S. R., and Tucker, C. S., 2018, "Human Validation of Computer vs Human Generated Design Sketches," ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (DETC2018-85698), Quebec City, Canada.
- [35] Chen, Y., Tu, S., Yi, Y., and Xu, L., 2017, "Sketch-pix2seq: A Model to Generate Sketches of Multiple Categories," arXiv preprint arXiv1709.04121.
- [36] Kaelbling, L. P., Littman, M. L., and Moore, A. W., 1996, "Reinforcement Learning: A Survey," Int. J. Artif. Intell. Res., 4, pp. 237–285.
- [37] Sutton, R. S., and Barto, A. G., 1998, Reinforcement Learning: An Introduction, MIT Press, Cambridge, MA.
- [38] Xu, X., Zuo, L., and Huang, Z., 2014, "Reinforcement Learning Algorithms With Function Approximation: Recent Advances and Applications," Inf. Sci., 261, pp. 1–31.
- [39] Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M., 2013, "Playing Atari With Deep Reinforcement Learning," arXiv Preprint arXiv:1312.5602.
- [40] Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., and Hassabis, D., 2017, "Mastering the Game of Go Without Human Knowledge," Nature, 550(7676), pp. 354–359.
- [41] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D., 2015, "Human-Level Control Through Deep Reinforcement Learning," Nature, 518(7540), pp. 529–533.
- [42] Johnson, M., Hofmann, K., Hutton, T., and Bignell, D., 2016, "The Malmo Platform for Artificial Intelligence Experimentation," IJCAI International Joint Conference on Artificial Intelligence, New York, NY, pp. 4246–4247.
- [43] Todorov, E., Erez, T., and Tassa, Y., 2012, "MuJoCo: A Physics Engine for Model-Based Control," IEEE International Conference on Intelligent Robots and Systems, Vilamoura, Algarve Portugal, pp. 5026–5033.
- [44] Juliani, A., Berges, V.-P., Vckay, E., Gao, Y., Henry, H., Mattar, M., and Lange, D., 2018, "Unity: A General Platform for Intelligent Agents," arXiv preprint arXiv:1809.02627.
- [45] Wang, P., Rowe, J., Min, W., Mott, B., and Lester, J., 2017, "Interactive Narrative Personalization With Deep Reinforcement Learning," IJCAI International Joint Conference on Artificial Intelligence, Melbourne, Australia, pp. 3852–3858.
- [46] Rowe, J., Smith, A., Pokorny, R., Mott, B., and Lester, J., 2018, "Toward Automated Scenario Generation With Deep Reinforcement Learning in GIFT," Proceedings of the Sixth Annual GIFT Users Symposium, Orlando, FL, pp. 65–74.
- [47] Hosseini-Nasab, H., Fereidouni, S., Fatemi Ghomi, S. M. T., and Fakhrzad, M. B., 2018, "Classification of Facility Layout Problems: A Review Study," Int. J. Adv. Manuf. Technol., 94(1-4), pp. 957-977.
- [48] Pardalos, P. M., Du, D. Z., and Graham, R. L., 2013, Handbook of Combinatorial Optimization, Springer, New York.

- [49] Drira, A., Pierreval, H., and Hajri-Gabouj, S., 2007, "Facility Layout Problems: A Survey," Annu. Rev. Control, 31(2), pp. 255–267.
- [50] Lotfi, N., and Acan, A., 2015, "Learning-Based Multi-Agent System for Solving Combinatorial Optimization Problems: A New Architecture," Proceedings of the 10th International Conference Hybrid Artificial Intelligent Systems, Bilbao, Spain, pp. 319–332.
- [51] Martin, S., Ouelhadj, D., Beullens, P., Ozcan, E., Juan, A. A., and Burke, E. K., 2016, "A Multi-Agent Based Cooperative Approach to Scheduling and Routing," Eur. J. Oper. Res., 254(1), pp. 169–178.
- [52] Samma, H., Lim, C. P., and Mohamad Saleh, J., 2016, "A New Reinforcement Learning-Based Memetic Particle Swarm Optimizer," Appl. Soft Comput. J., 43, pp. 276–297.
- [53] Silva, M. A. L., De Souza, S. R., Souza, M. J. F., and De Oliveira, S. M., 2015, "A Multi-Agent Metaheuristic Optimization Framework With Cooperation," Proceedings—2015 Brazilian Conference on Intelligent Systems (BRACIS), Natal, Brazil, pp. 104–109.
- [54] Aydin, M. E., and Öztemel, E., 2000, "Dynamic Job-Shop Scheduling Using Reinforcement Learning Agents," Rob. Autom. Syst., 33(2-3), pp. 169–178.
  [55] Wang, Y. C., and Usher, J. M., 2005, "Application of Reinforcement Learning
- [55] Wang, Y. C., and Usher, J. M., 2005, "Application of Reinforcement Learning for Agent-Based Production Scheduling," Eng. Appl. Artif. Intell., 18(1), pp. 73–82.
- [56] Silva, M. A. L., de Souza, S. R., Souza, M. J. F., and Bazzan, A. L. C., 2019, "A Reinforcement Learning-Based Multi-Agent Framework Applied for Solving Routing and Scheduling Problems," Expert Syst. Appl., 131, pp. 148–171.
- Routing and Scheduling Problems," Expert Syst. Appl., 131, pp. 148–171.
  [57] Shahrabi, J., Adibi, M. A., and Mahootchi, M., 2017, "A Reinforcement Learning Approach to Parameter Estimation in Dynamic Job Shop Scheduling," Comput. Ind. Eng., 110, pp. 75–82.
- [58] Nazari, M., Oroojlooy, A., Snyder, L. V., and Takáč, M., 2018, "Reinforcement Learning for Solving the Vehicle Routing Problem," 32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montréal, Canada, pp. 9839– 9849
- [59] Govindaiah, S., and Petty, M. D., 2019, "Applying Reinforcement Learning to Plan Manufacturing Material Handling Part 2: Experimentation and Results," ACMSE 2019—Proceedings of the 2019 ACM Southeast Conference, Kennesaw, GA, pp. 16–23.

- [60] Govindaiah, S., and Petty, M. D., 2019, "Applying Reinforcement Learning to Plan Manufacturing Material Handling Part 1: Background and Formal Problem Specification," ACMSE 2019—Proceedings of the 2019 ACM Southeast Conference, Kennesaw, GA, pp. 168–171.
- [61] Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O., 2017, "Proximal Policy Optimization Algorithms," arXiv Preprint arXiv:1707.06347.
- [62] Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P., 2015, "Trust Region Policy Optimization," International Conference on Machine Learning, pp. 1889–1897.
- [63] Lopez, C., Ashour, O., and Tucker, C., 2019, "An Introduction to CLICK: Leveraging Virtual Reality to Integrate the Industrial Engineering Curriculum," ASEE Annual Conference & Exposition, Tampa, FL, pp. 1–12.
- [64] Terpenny, J. P., Harmonosky, C. M., Lehtihet, E., Prabhu, V. V., Freivalds, A., Joshi, E. M., Ventura, J. A., 2018, "Product-Based Learning: Bundling Goods and Services for an Integrated Context-Rich Industrial Engineering Curriculum," Annual Conference of the American Society for Engineering Education (ASEE), Salt Lake City, UT.
- [65] W. G., and Pope, C., 2011, "Unity\_Game\_Engine," UNITY GAME ENGINE Overy.
- [66] Petridis, P., Dunwell, I., De Freitas, S., and Panzoli, D., 2010, "An Engine Selection Methodology for High Fidelity Serious Games," 2nd International Conference on Games and Virtual Worlds for Serious Applications, Braga, Portugal, pp. 27–34.
- [67] Alsubaie, A., Alaithan, M., Boubaid, M., and Zaman, N., 2018, "Making Learning Fun: Educational Concepts & Logics Through Game," International Conference on Advanced Communication Technology (ICACT), Gang'weondo, South Korea, pp. 454–459.
- [68] Cunningham, J., and Tucker, C. S., 2018, "A Validation Neural Network (VNN) Metamodel for Predicting the Performance of Deep Generative Designs," Proceedings of ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (DETC2018-86299), Quebec City, Quebec, Canada.
- [69] Burda, Y., Edwards, H., Pathak, D., Storkey, A., Darrell, T., and Efros, A. A., 2018, "Large-Scale Study of Curiosity-Driven Learning," arXiv preprint, arXiv: 1808.04355.