

# Online Game Level Generation from Music

Ziqi Wang<sup>\*†</sup>, Jialin Liu<sup>†\*</sup>

<sup>\*</sup> *Research Institute of Trustworthy Autonomous System  
Southern University of Science and Technology (SUSTech)*

<sup>†</sup> *Guangdong Provincial Key Laboratory of Brain-inspired Intelligent Computation  
Department of Computer Science and Engineering  
Southern University of Science and Technology (SUSTech)  
Shenzhen, China*

**Abstract**—Game consists of multiple types of content, while the harmony of different content types play an essential role in game design. However, most works on procedural content generation consider only one type of content at a time. In this paper, we propose and formulate online level generation from music, in a way of matching a level feature to a music feature in real-time, while adapting to players’ play speed. A generic framework named online player-adaptive procedural content generation via reinforcement learning, OPARL for short, is built upon the experience-driven reinforcement learning and controllable reinforcement learning, to enable online level generation from music. Furthermore, a novel control policy based on local search and k-nearest neighbours is proposed and integrated into OPARL to control the level generator considering the play data collected online. Results of simulation-based experiments show that our implementation of OPARL is competent to generate playable levels with difficulty degree matched to the “energy” dynamic of music for different artificial players in an online fashion.

**Index Terms**—Procedural content generation, online level generation, player-adaptive, EDPCG, EDRL

## I. INTRODUCTION

Human perception is multi-modal. Digital games, as an emerging creation field, lies in the intersection of multiple types of content that meet different aspects of human perception, has the ability of expressing stories, emotions or aesthetics, and satisfying human’s natural entertainment demand [1], [2]. A successful game should guarantee the harmony of different types of content. Procedural content generation (PCG) [3]–[6], aiming at the automated or mixed-initiative creation of game contents, such as levels, maps, musics, rules and narratives, has shown its potential to reduce game development costs, augment the creativity of individual human creators, and provide personalised game contents [2], [6]. There have been extensive researches in generating a single type of game content [2], [6], while only a few works considered generating one type of content driven by another [7]–[9]. Procedurally

This work was supported by the Guangdong Provincial Key Laboratory (Grant No. 2020B121201001), the Program for Guangdong Introducing Innovative and Entrepreneurial Teams (Grant No. 2017ZT07X386), the Shenzhen Science and Technology Program (Grant No. KQTD2016112514355531), the Shenzhen Fundamental Research Program (Grant No. JCYJ20190809121403553), the National Natural Science Foundation of China (Grant No. 61906083), the Research Institute of Trustworthy Autonomous Systems (RITAS) and the SUSTech Undergraduate Teaching Quality and Reform Project (Grant No. SJZLGC202101).

Corresponding author: Jialin Liu (liujl@sustech.edu.cn).

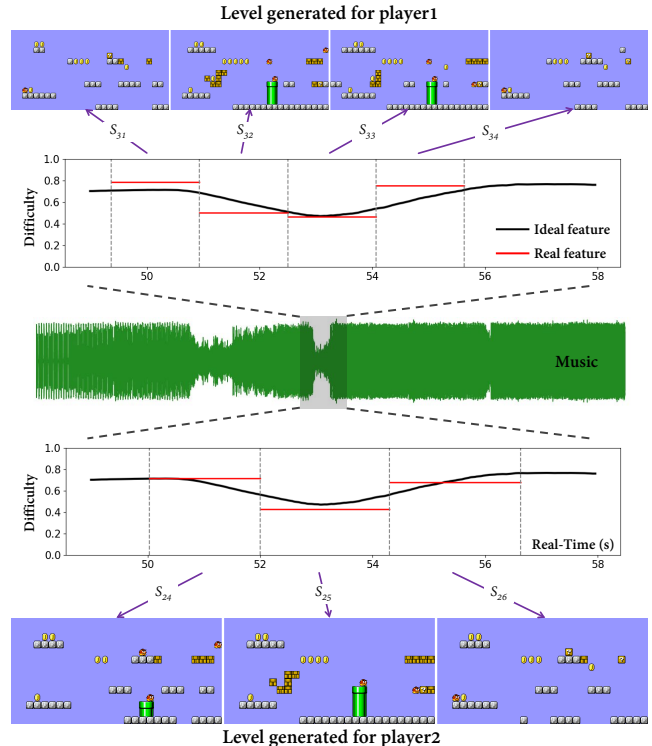


Fig. 1: Illustration of online level generation from music. Widths of interval between grey dashed lines correspond to the time consumed to play through a segment. This figure is plotted based on the data of simulation-based experiments of this paper (detailed in Section V). The segment-wise difficulty degrees are closed to the ideal difficulty curve in black colour, which is converted from the input music’s energy. Moreover, the level generation is adapted to the specific player’s play speed in real-time.

generating a complete game of several content types is an emerging topic [6], [10]. For instance, *AudioInSpace* [10] generated game rules, visuals and audio in a mix-initiative way. The aforementioned works [1], [7]–[10] mainly focused on puzzle games, shooting games and rhythm games.

In this work, we focus on platformer games, in which levels and music together affect the player experience. We investigate into level generation from music, ideally to be achieved in an

online fashion for some games, assuming that the play experience should be consistent with the current music. For example, if the music in a platformer game is intense or nervous, the appeared level segments should be more difficult to match the atmosphere created by music. Some commercial games, such as *Dance Dance Revolution* (Konami, 1998), *Guitar Hero* (RedOctane, 2005) and *Muse Dash* (PeroPeroGames, 2018), directly force the gameplay to be consistent with music by designing specific game rules. However, for most games of other genres, it’s hard to generate levels from music because the rhythm of playing a level depends on the player’s play speed and thus not fixed, which will result in the failure of matching a level to music in an offline or player-unaware way. Fig. 1 illustrates why the play speed affects the online generation of level from music. The level segments on top and bottom are played by different artificial players at the same time window but have different length as the players played the levels in different speed. Moreover, the generation system should determine the difficulty degree of the next level segment in real-time to make sure that the segment will be played at the best-matched time slice. Therefore, online and player-adaptive level generation from music is desired.

In this paper, a framework for online level generation from music named online player-adaptive procedural content generation via reinforcement learning, OPARL for short, is proposed. OPARL follows a controller-generator architecture. The generator takes the current level segment and a music feature as input and outputs a new level segment with feature matched to the input at each iteration. When OPARL generates a level segment from music, the generator is controlled by a novel policy named local search with k-nearest neighbours based estimation (LS-KNN) which takes historical play data into account to produce the targeted feature value that minimises an inner error respect to an ideal feature sequence derived from the given music’s temporal feature sequence (detailed in section III). The controllable generator is extended from the experience-driven procedural content generation via reinforcement learning (EDRL) [11] framework by leveraging controllable reinforcement learning [12], [13]. The proposed OPARL framework is implemented and verified on the benchmark game for level generation, *Super Mario Bros.* (SMB). Experimental results show that the resulted system can generate playable SMB levels that are consistent with the given music<sup>1</sup>.

Our framework requires no expert knowledge except for the content representation, a few training examples for training a generative adversarial network (GAN) [14], [15] and the CNet-assisted repairer [16] for determining and repairing broken pipes. Although our proposed approach is verified in generating SMB levels from music, applying it to other platformer games, such as *Megaman* (CAPCOM, 1987), *Electronic Super Joy* (Michael Todd, 2014) and *Celeste* (Matt Makes Games Inc., 2018), is straightforward.

<sup>1</sup>Code, experimental data and demo of this paper are available on GitHub: <https://github.com/PneuC/OPARL>.

The remainder of this paper is organised as follows. Section II presents some related work. Section III formulates the problem of online level generation from music and addresses its challenges. Then, our OPARL framework and its technical details are presented in Section IV. In Section V, the effectiveness and robustness of OPARL to different players and music are verified through its implementation for SMB and simulation-based experiments. Section VI concludes and discusses some future directions.

## II. BACKGROUND

This section discusses related work that involves multiple types of content and online player-adaptive approach.

### A. PCG that Involves Multiple Content Types

Liapis *et al.* proposed the concept of *orchestrating game generation*, which aims at generating different types of game content jointly and harmoniously [1]. One representative orchestrating game generation system, *AudioInSpace* [10], generates game rules, visuals and audio in a mix-initiative way. Plans and Morelli proposed an experience-driven generator to generate music that reacts to the “excitement” of the game using search-based algorithms [7]. Naushad and Muhammad introduced a conditional music generation framework to enable adaptive music generation for games [8]. Engels *et al.* developed an hierarchical Markov model-based music generation system to produce music pieces in real-time [9].

The aforementioned works mainly focus on generating *cosmetic* content from *functional* content [5], while there are also research works that explored the reversed way. For instance, some works concentrated on learning to generate rhythms game charts from music via supervised learning [17], [18]. Jordan *et al.* introduced a mobile game named *BeatTheBeat* which applied self-organising maps method to create game levels that match some music features [19]. In the work of [20], a mixed-initiative PCG system is presented by Karavolos *et al.* to generate game levels from mission or space provided by human designers. Atmaja *et al.* proposed a top-down PCG framework to “translate” platformer games from storyline [21]. To our best knowledge, no work has ever generated platformer game levels from music in real-time.

### B. Online Player-Adaptive Level Generation

A popular research topic related to online player-adaptive level generation is dynamic difficult adjustment (DDA), aiming at adjusting levels’ difficulty degree considering the abilities or skills of players for desired player experience or aesthetic goal [22]. Shi and Chen proposed a DDA policy based on Thompson sampling, and embedded it into an online level generation framework [23]. Stammer *et al.* applied a conditional player experience model considering different play styles to generate player-adaptive *Spelunky* levels with DDA [24].

Some works built on experience-driven procedural content generation [25]. For instance, Shaker *et al.* applied player modelling and grammar evolution to generate online levels that optimise player experience [26]. Blom *et al.* generated online

personalised SMB levels with facial expression recognition [27]. Shu *et al.* introduced the EDRL framework to enable real-time level generation with experience-driven reward functions as content quality measurements [11].

### III. MUSIC-DRIVEN ONLINE LEVEL FEATURE CONTROL

To generate level segments that are consistent to a given piece of music in real-time, we consider ensuring the consistency or harmony by matching a feature of level segments to a feature of the given piece of music, referred to as music-driven online level feature control in this paper. The problem of finding online the optimal value of level feature that matches the given music is formulated in Section III-A). Then, the challenges of achieving online music-driven level feature control are discussed in Section III-B.

#### A. Problem Formulation

Given a piece of music, an ideal feature sequence of a level can be derived in some fine-grained time unit, denoted as  $\mathcal{F}^* = (f_1^*, \dots, f_t^*, \dots)$ . Let  $\hat{f}_i$  and  $\delta_i$  denote the target feature produced by the controller at the  $i$ th iteration and the duration of playing the  $i$ th segment,  $S_i$ , respectively. The objective is to minimise the error defined as follows:

$$\varepsilon_{inner} = \frac{1}{T(f_{max} - f_{min})} \sum_{i=1}^N \left( \sum_{t=b_i}^{b_i+\delta_i} |f_t^* - \hat{f}_i| \right), \quad (1)$$

where  $N$  is the number of segments and  $b_i = \sum_{k=0}^{i-1} \delta_k$  indicates the starting time of playing the level segment  $S_i$  with  $\delta_0 = 0$ .  $f_{max}$  and  $f_{min}$  are the upper bound and lower bound of the level feature, respectively.  $T = \sum_{i=1}^N \delta_i$  is the total time spent to play through the whole level. The decision space of online music-driven level feature control is  $[f_{min}, f_{max}]$ .

This error  $\varepsilon_{inner}$  is called an *inner error* because Eq. (1) evaluates the distance between the targeted feature values and the ideal feature sequence, and cannot be eliminated as  $\delta_i$  is always larger than the time unit. In a controller-generator architecture, there can be an additional *outer error* between the targeted feature value and the one of an actually generated segment. Hence, the *overall error* does not equal to *inner error*. An *outer error*,  $\varepsilon_{outer}$ , and an *overall error*,  $\varepsilon_{all}$ , are formulated in Eqs. (2) and (3), respectively.

$$\varepsilon_{outer} = \frac{1}{T(f_{max} - f_{min})} \sum_{i=1}^N \delta_i |\hat{f}_i - f_i|. \quad (2)$$

$$\varepsilon_{all} = \frac{1}{T(f_{max} - f_{min})} \sum_{i=1}^N \left( \sum_{t=b_i}^{b_i+\delta_i} |f_t^* - f_i| \right). \quad (3)$$

The notation  $f_i$  in Eq (2) denotes the real feature value of  $S_i$ . As shown in Eqs. (1), (2) and (3), those errors are normalised linearly to the range of  $[0, 1]$ . Fig. 2 illustrates the relationship between those errors.

Although only one level feature is considered as a case study in this work, our formulation can be easily extended to the case of using multiple level features.




Fig. 2: Illustration of the inner error, outer error and overall error formulated in Eqs. (1), (2) and (3), respectively.

#### B. Challenges of Music-driven Online Level Feature Control

There are at least two challenges of music-driven online level feature control, the uncertainty of play duration and the dilemma of granularity. The time consumed to play through a segment depends on the player’s skill and play style. Therefore, generating level segment from music should be achieved in a player-adaptive way. Online level generation usually generates a level segment by segment. Using segments of smaller size is expected to control features more accurately. However, online level generation also requires high generation speed. Using smaller segments will lead to a lower generation speed due to the higher frequency of making control decision. Moreover, extracting features from very small segments does not always make sense. Generating small yet reasonable segments in real-time is not trivial. To overcome the above challenges, a music-driven online level feature controller is designed and detailed in Section IV-C.

### IV. ONLINE PLAYER-ADAPTIVE PROCEDURAL CONTENT GENERATION VIA REINFORCEMENT LEARNING

We propose a framework named online player-adaptive procedural content generation via reinforcement learning (OPARL) for online level generation from music. An overview of OPARL is given in Fig. 3. OPARL is composed of a feature controller and a segment generator. At each iteration, the controller receives players’ play data on the most recent segment and determines the control signal, which directs the generator to generate appropriate segments for a specific player. The generator takes a number of previous segments and the control signal as input, and then outputs a new segment. This section explains the framework of OPARL, its implementation and parameter setting used in the experimental studies. More technical details are available in the released project<sup>1</sup>.

#### A. Controllable EDRL with Archive

A controllable experience-driven reinforcement learning with archive (CEDRL-A) architecture, extended from the EDRL framework [11], is designed as the generator of OPARL. It models the task of online level generation as a Markov decision process so that a *designer* (known as “agent” in [11]) can be trained to generate rapidly level segments of high quality. A *booster* (known as “generator” in [11]) based on GAN is applied to tackle the high-dimensionality of designer’s action space. The booster maps a low-dimensional

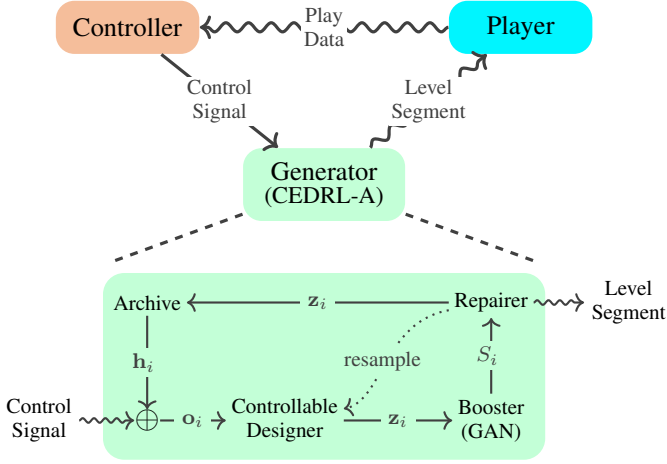


Fig. 3: Overview of OPARL.  $\mathbf{h}_i$ ,  $\mathbf{o}_i$  and  $\mathbf{z}_i$  refer to the historical latent vectors, observation of designer and latent vector at the  $i$ th iteration, respectively. *Archive* is a ring queue of latent vectors with capacity  $m_G$ .

vector (e.g., latent vector of GAN) to a high-dimensional level segment representation. The observation and action of the designer are the latent vectors of the most recently generated segment and the segment to be generated, respectively. The action space of CEDRL-A is the latent space of the booster.

Different to EDRL, the designer in CEDRL-A takes additional control signal that describes the desired feature value of the next segment as input, and outputs a segment with a feature value closed to what the input control signal described. This ability of generating segments with a desired feature value is achieved by employing a *controllability* reward while training the generator. Moreover, an *archive*  $X_G = \{\mathbf{z}_j \mid j = i - m_G, \dots, i - 1\}$  of  $m_G$  previous latent vectors is introduced in CEDRL-A to guarantee the Markov property. Since some reward functions can depend on more than one previous segments, it is necessary to make sure that  $m_G$  is larger than or equal to the maximum number of previous segments to be used in any reward function. Let  $\mathbf{h}_i = \mathbf{z}_{i-m_G} \oplus \dots \oplus \mathbf{z}_{i-1}$ <sup>2</sup> represent the concatenated historical latent vectors that a controller can observe and  $\mathbf{g}_i$  represent the vector of the control signal, the controllable designer takes  $\mathbf{o}_i = \mathbf{h}_i \oplus \mathbf{g}_i$  as its observation at the  $i$ th iteration.

Similar to EDRL, CEDRL-A applies a CNet-assisted evolutionary repairer to repair the broken pipes appeared in the generate levels [16] and a resampling strategy to ensure the playability [11], referred to as *Repairer* in Fig. 3.

## B. Training Controllable Designer

1) *Reward function*: To make the designer generate segments with feature values closed to the desired ones at each iteration, a *controllability* reward is introduced during designer

training. Let  $f(\cdot)$  be the function of computing level feature, the controllability reward is formulated as

$$C(S_i) = 1 - \frac{|f(S_i) - \hat{f}_i|}{f_{max} - f_{min}}.$$

In the implementation of OPARL in this paper, the difficulty degree of a level segments is considered as the feature of a level segment, quantified as the summation of the number of enemies and the number of empty ground tiles divided by the width of segment, as formulated as

$$\text{difficulty} = \frac{\#\text{enemies} + \#\text{empty\_grounds}}{\text{width}}. \quad (4)$$

The lower bound and upper bound of difficulty degree are set as  $f_{min} = 0$  and  $f_{max} = 1$ , respectively.

As the targeted feature (difficulty degree in the implementation) directly influences the next segment, the method of sampling features can affect other rewards in some unpredictable way. In some primary attempts, we observe that the difficulty degree usually does not change fast in the online generation phase. Therefore, the feature is sampled uniformly in  $[f_{min}, f_{max}]$ , and then a Gaussian mutation is applied to the current feature value to get the next feature value recursively. Formally, the feature values sampled during a training epoch are  $\hat{f}_1 \sim U(f_{min}, f_{max})$  and  $\hat{f}_i = \hat{f}_{i-1} + \xi$  with  $\xi \sim N(0, \sigma^2)$ ,  $\forall i > 1$ . In this work,  $\sigma$  is arbitrarily set as 0.05.

We also employ *fun* and *playability* introduced in the work of [11] as additional reward terms so that the generated levels are fun and playable. The playability of segment  $S_i$  is checked by simulating  $S_{i-1} + S_i$  with an  $A^*$  agent that won the 2009 Mario AI Competition [28]. In our work, the *playability* is set as 0 if the newly generated segment is playable, otherwise  $-1$ .

2) *Training Designer with Soft Actor-critic*: Among the reward terms, *controllability* does not depend on previous segments, and *playability* depends on no more than one previous segment. However, in our setting, computing *fun* requires 2 previous segments ( $S_{i-1}$  and  $S_i$ ). Thus, the capacity of the generator's archive  $m_G$  is 2. The designer uses a multi-layer perceptron (MLP) model. The booster uses a latent space of  $\mathbb{R}^{20}$ . The targeted difficulty degree is duplicated by 12 times to increase the number of connections in the MLP on the inputted targeted difficulty. As a result, the control signal  $\mathbf{g}_i$  has a dimensionality of 12 and the designer's observation  $\mathbf{o}_i$  is a 52-dimensional vector. The action space is  $[-1, 1]^{20}$ .

We implement a parallel-environment version of soft actor-critic (SAC) [29], [30] as an OpenAI Gym interface [31] with 50 synchronous sub-environments to train the designer for 1 million time steps in total. The actor and critics use MLP model with 3 hidden layers of 256 neurons. Our implementation of SAC updates the models 10 times using randomly sampled batches of size 384 every time 100 transitions are collected. The automating entropy adjustment [30] is used with a targeted entropy  $\mathcal{H} = -20$ , as recommended in [30]. The smoothing coefficient  $\tau$  is set as 0.005. The discounted factor  $\gamma$  greatly affects the performance of reinforcement learning algorithms. Empirically, a smaller  $\gamma$  (comparing with  $\gamma = 0.99$

<sup>2</sup> $\oplus$  represents vector concatenating.

in many baselines) is better. In all the experiments reported in this paper,  $\gamma = 0.7$ . Each epoch terminates after 50 segments are generated. For the last transition, the target Q-value is computed as  $r + \gamma Q(s', \tilde{a}')$  rather than  $r$  since the training aims at endless online generation.

### C. Online Player-Adaptive Level Generation with Controller

Given a piece of music, our music-driven feature controller will first compute the ideal feature sequence  $\mathcal{F}^*$  during initialisation, and then store it. In addition to  $\mathcal{F}^*$ , the control policy keeps an archive  $X_C = \{\langle \hat{f}_j, \delta_j \rangle \mid j = i - m_C - 1, \dots, i - 2\}$  to keep at most  $m_C$  entries of the previous segments. Each entry is composed of the targeted feature value and play duration of a segment. The last entry of  $X_C$  is about the  $(i - 2)$ th segment at the  $i$ th iteration because  $\delta_{i-1}$  is unknown when determining  $\hat{f}_i$ .

---

**Algorithm 1** LS-KNN. In the experiments,  $\#\text{trial} = 50$ ,  $k = 5$  and  $\sigma_c = 0.02$ .

---

**Require:**  $\hat{f}_{i-1}, b_{i-1}, \mathcal{F}^*, X_C$   $\triangleright$  Inputs  
**Require:**  $\#\text{trials}, k, \sigma_c$   $\triangleright$  Hyper-parameters  
**Ensure:**  $\hat{f}_i$   
1:  $\tilde{b}_i \leftarrow b_{i-1} + \text{Estimate}(\hat{f}_{i-1}; X_C)$   $\triangleright$  Estimate beginning time  
2:  $\hat{f}_i \leftarrow \hat{f}_{i-1}$   
3:  $\tilde{\delta}_i \leftarrow \text{Estimate}(\hat{f}_i; X_C)$   
4:  $\tilde{\varepsilon} \leftarrow \text{Evaluate}(\hat{f}_i; \tilde{b}_i, \tilde{\delta}_i, \mathcal{F}^*)$   $\triangleright$  Least estimated error found  
5: **repeat**  
6:  $\hat{f}'_i \leftarrow \hat{f}_i + \xi, \quad \xi \sim N(0, \sigma_c^2)$   $\triangleright$  Do local search  
7:  $\tilde{\delta}'_i \leftarrow \text{Estimate}(\hat{f}'_i; X_C)$   
8:  $\tilde{\varepsilon}' \leftarrow \text{Evaluate}(\hat{f}'_i; \tilde{b}_i, \tilde{\delta}'_i, \mathcal{F}^*)$   
9: **if**  $\tilde{\varepsilon}' < \tilde{\varepsilon}$  **then**  
10:  $\hat{f}_i \leftarrow \hat{f}'_i$   
11:  $\tilde{\varepsilon} \leftarrow \tilde{\varepsilon}'$   
12: **end if**  
13: **until** has looped for  $\#\text{trials}$  times  
14: **return**  $\hat{f}_i$

---

A simple algorithm named local search with KNN-based estimation (LS-KNN, Algorithm 1) is designed as the online control policy for determining targeted feature values. When determining a targeted feature value for the  $i$ th segment, LS-KNN executes local search for  $\#\text{trials}$  generations starting from the last one  $\hat{f}_{i-1}$ , and picks up the best value of  $\hat{f}_i$  found according to the estimated individual *inner error* using KNN-based prediction of play duration for the  $(i - 1)$ th and  $i$ th segments. In our case, an individual is mutated by adding a Gaussian noise with a standard deviation of  $\sigma_c$ . The play duration is estimated with:

$$\text{Estimate}(\hat{f}; X_C) = \frac{1}{k} \sum_{j \in J} \delta_j,$$

where  $J$  is the set of  $k$  nearest neighbours in terms of  $|\hat{f}_j - \hat{f}|$  within  $X_C$ . With  $\text{Estimate}(\hat{f}; X_C)$ , we can further estimate

the time that the  $i$ th segment  $\tilde{b}_i$  starts to be played. With  $\tilde{\delta}_i$ , the estimated play duration respected to  $\hat{f}_i$ , a feature can be evaluated using:

$$\text{Evaluate}(\hat{f}_i; \tilde{b}_i, \tilde{\delta}_i, \mathcal{F}^*) = \frac{1}{\tilde{\delta}_i} \sum_{t=\tilde{b}}^{\tilde{b}+\tilde{\delta}_i} |f_t^* - \hat{f}_i|.$$

Algorithm 1 details the implementation of the control policy. LS-KNN is proposed with two assumptions: (i) an ideal feature sequence won't change fast, thus the "optimal"  $\hat{f}_i$  should not be far from  $\hat{f}_{i-1}$ ; (ii) the difficulty implies the play duration, which is true for many platformer games. The former is somehow generic in online level generation because fast changes of features may be harmful to the coherence of levels. It is the reason of starting the local search from  $\hat{f}_{i-1}$  in Algorithm 1. The latter motivates the estimation of play duration according to the records organised by  $\langle \hat{f}_j, \delta_j \rangle$ .

LS-KNN is easy to implement and can always be used directly without training or other preparation. Furthermore, though not explicit, LS-KNN is well player-adaptive since the KNN-based estimation is applied based on the specific player's play data collected online. Our implementation with LS-KNN policy achieves significant performance in the simulation-based experiments detailed in Section V.

### D. Implementation of Booster

Our booster is a variation of MarioGAN [15]. Differ from the original neural network architecture, our GAN model uses fractional-convolutional layers [32] with kernel size, stride and padding of  $\langle (4, 4), (1, 1), 0 \rangle$ ,  $\langle (3, 3), (2, 2), 1 \rangle$ ,  $\langle (4, 4), (2, 2), 1 \rangle$  and  $\langle (3, 4), (1, 2), 1 \rangle$ , respectively, to directly obtain an output of size  $14 \times 28$  without clipping. Besides, latent vectors of length 20 are used. The generator and discriminator are trained for 5 times and 1 time at each iteration, respectively, on the 13 human-designed levels without bullet bills provided in the Video Game Level Corpus [33].

## V. EXPERIMENTAL STUDY

To evaluate the effectiveness of our approach and implemented algorithms in optimising different objectives, the designer is trained with all the possible combinations of the three reward terms presented in Section IV-B1 with the same weights, and evaluated with training environments and online generation environments, respectively.

To test the robustness of our method, five different agents in the 2009 Mario AI Competition [28], namely Baumgarten's (the aforementioned  $A^*$  agent), Sloane's, Hartmann's, Polikarpov's and Schumann's agents, are used as the simulated player, and two different pieces of music, *Ginseng*<sup>3</sup> (EnV, 2014) and *Farewell*<sup>4</sup> (Raine, 2019), are used. Fig. 4. shows the five agents' play duration on each segments of an level generated online by OPARL using a designer trained with the summation of *controllability*, *fun* and *playability*. Those agents

<sup>3</sup>From the original sound track of commercial platformer game *Electronic Super Joy: Groove City* (Michael Todd, 2014).

<sup>4</sup>From the original sound track of commercial platformer game *Celeste* (Matt Makes Games Inc., 2018).




Fig. 4: Cumulative play duration of each agent on levels generated from the same initial segment, averaged over 30 independent trails. Shadows indicate the standard deviation.

actually play levels with different speeds. Our experiments are simulated based on the Mario-AI-Framework<sup>5</sup>.

Root-mean-square energy features of music are extracted throughout time and mapped to difficulty degrees as follows. First, the degrees are re-scaled by taking logarithm based on 10, then clipped within  $[-2.5, 0]$  and mapped into  $[0, 1]$  linearly, finally, smoothed through computing mean values in an 1-stride sliding window of size 100. The music feature extraction is done through *Librosa* library [34], with a default time unit of 0.02322s. Parts of the resulted ideal difficulty sequences are illustrated in Fig. 5. Demos of levels generated from different music are available in the released project<sup>1</sup>.




Fig. 5: Parts of converted ideal difficulty sequences extracted from the music used in our experiments.

#### A. Evaluation of Effectiveness

The trained designers are evaluated with three metrics,  $\sqrt{-F}$ , the mean distance of  $D(S_i)$  out of the range  $[l, u]$ ,  $-P$ , the negative number of *playability* to indicate the ratio of unplayable segments, and  $1 - C$  to indicate the mean error between level feature and targeted feature value in training environment. Those metrics are averaged over all iterations (time steps) and normalised in a similar scale.

To further evaluate the performance of online generation, the designer is tested using Baumgarten’s agent as simulated player and *Ginseng* as input music. The values of  $\sqrt{-F}$ ,  $-P$ ,  $\epsilon_{inner}$ ,  $\epsilon_{outer}$ ,  $\epsilon_{all}$ , and an additional metric *Div*, which is used to measure the *diversity* of levels generated in different runs, are reported in Table I. *Div* is computed as the mean ratio of tiles that are different in pairs of levels generated in different runs. The values of  $\sqrt{-F}$  and  $-P$  in the online evaluation may be different to the ones in the offline evaluation due to resampling and different targeted feature values.

Table I shows the experimental results. The designers trained with *controllability* generally achieve very low *overall error* in the online generation tests. The main source of *overall error* is the *outer error*, i.e., the error between the targeted feature value produced by controller and the feature of actually generated segment. The value of  $1 - C$  closed to  $\epsilon_{outer}$  means that our method of sampling targeted features is effective. The designer trained with only *fun* reward achieves a great score on  $\sqrt{-F}$ . However, when *controllability* is employed, the score of *fun* deteriorates a lot. This phenomenon indicates that the objective of *fun* and *controllability* conflict. Moreover, the *controllability* deteriorates less comparing with the designer trained with *controllability* only. It is probably because the reward of *fun* uses a quadratic form while *controllability* uses a linear form. That means designer finds it better to optimise *controllability* to get a higher summation of reward terms.

All the designers trained with *playability* well ensure the playability of generated levels, while designers generally assure better playability with the help of resampling. A merit attention finding is that the designer trained with  $F$  and the designer trained with  $C + F + P$  do not get notable better  $P$  value in the online generation tests. A possible reason is that those designers lack of randomness when taking actions. That means if they generate an unplayable segment, no matter how many times the re-generation is executed, they will still generate unplayable segments. A future work is finding out why the phenomenon only appears on those two designers.

As a conclusion, our implemented framework optimises the reward functions effectively. The designer trained with  $C + F + P$  balances different objectives and can be a good choice for online level generation from music.

#### B. Evaluation of Robustness


Fig. 6 plots the *overall error*, *fun* and *diversity* evaluated on the designer trained with  $C + F + P$  as reward for the five agents and two different pieces of music. The *overall error* and *fun* are plotted as  $1 - \epsilon_{all}$  and  $1 - \sqrt{-F}$  for better intelligibility. According to Figs. 6(a) and 6(b), our method achieves very similar and high performances of *overall error* and *fun*, and is robust for players with different play speed. The diversity scores of levels generated for different players and musics vary significantly. According to Figs. 4 and 5, the diversity of levels generated by OPARL may be positively correlated with the fluctuation degree of music and the variance of play duration.

Fig. 7 presents segments captured from the levels generated with the same  $S_0$  (i.e., initial segment) for different agents and different music. Fig. 1 uses an illustration to explain how OPARL generates different levels for different players. It is shown in Fig. 7 that levels generated from *Ginseng* are generally harder than the ones generated from *Farewell*, as the ideal difficulty sequence derived from *Ginseng* is generally larger than the one from *Farewell* (cf. Fig. 5). The levels generated for different agents with the same music are similar. It may be explained by using the same starting segment to


<sup>5</sup><https://github.com/amidos2006/Mario-AI-Framework>

TABLE I: Evaluation of designers trained with different reward functions. All values are averaged over 100 independent trials.  $Div$  is to be maximised, all the other metrics are to be minimised with strict lower bound of 0. Cells filled with  $-/-$  are meaningless. The best and worst results are highlighted with bold and italic, respectively.


Designer	Training Environment			Online Generation					
	$\sqrt{-F}(10^{-2})$	$-P(10^{-2})$	$1 - C(10^{-2})$	$\sqrt{-F}(10^{-2})$	$-P(10^{-2})$	$\epsilon_{inner}(10^{-2})$	$\epsilon_{outer}(10^{-2})$	$\epsilon_{all}(10^{-2})$	$Div$
$F$	<b>1.08 ± 0.34</b>	9.12 ± 4.03	-/-	<b>1.06 ± 0.35</b>	9.06 ± 4.22	-/-	-/-	43.6 ± 1.76	<b>0.066</b>
$P$	11.2 ± 1.81	<b>0.22 ± 0.63</b>	-/-	11.2 ± 1.90	<b>0.00 ± 0.00</b>	-/-	-/-	48.2 ± 1.01	<b>0.066</b>
$F+P$	2.06 ± 0.58	0.42 ± 0.82	-/-	2.29 ± 0.58	<b>0.00 ± 0.00</b>	-/-	-/-	49.9 ± 0.92	0.046
$C$	16.8 ± 3.27	46.7 ± 21.2	<b>1.90 ± 0.23</b>	19.2 ± 0.99	29.4 ± 5.49	1.16 ± 0.10	<b>1.59 ± 0.22</b>	<b>2.24 ± 0.19</b>	0.056
$C+F$	4.48 ± 2.13	15.3 ± 6.76	2.88 ± 0.49	1.83 ± 0.52	5.66 ± 3.12	0.89 ± 0.04	2.69 ± 0.26	2.95 ± 0.24	0.057
$C+P$	11.9 ± 3.41	0.40 ± 0.85	3.77 ± 0.70	17.5 ± 1.59	0.44 ± 0.96	0.89 ± 0.04	4.82 ± 0.73	4.99 ± 0.80	0.062
$C+F+P$	6.81 ± 1.31	0.40 ± 0.94	4.36 ± 0.56	7.86 ± 0.97	0.02 ± 0.20	<b>0.87 ± 0.03</b>	4.74 ± 0.54	4.93 ± 0.51	0.035



(a) Evaluation scores of  $1 - \epsilon_{all}$  for different simulated players.



(b) Evaluation scores of  $1 - \sqrt{-F}$  for different simulated players.



(c) Evaluation scores of  $Div$  for different simulated players.

Fig. 6: Values of  $fun$ ,  $controllability$ ,  $diversity$  evaluated on designers trained with different agents as player and different musics. Each value is averaged over 30 independent trials.

generate those levels. To summarise, our method can adapt well different players and is robust to different music.

## VI. CONCLUSION

This paper formulates the problem of online level generation from music, and proposes an online player-adaptive procedural content generation via reinforcement learning (OPARL) framework composed of a novel CEDRL-A generator and a novel LS-KNN controller to achieve online level generation from music. Experimental results show that the implementation of OPARL can generate in real-time SMB levels with segment-wise features closed to an ideal difficulty sequence derived from a piece of music. The resulted generation system can also guarantee the playability. The training algorithm implemented in this paper achieves considerable performance and can be used as a baseline in further studies. Moreover, our framework

is flexible since the controller and the generator are decoupled. The CEDRL-A generator in our framework can be integrated with other controllers like DDA controller for different aspects of player-adaptation.

In this paper, our proposed approaches are verified with simulation-based studies. One of the future work is conducting human tests. As another future work, new ways of mapping multiple features of both levels and music can be studied for the purpose of achieving better consistence between play experience and music.

## ACKNOWLEDGEMENT

The authors would like to thank the anonymous reviewers for their valuable comments.

## REFERENCES

- [1] A. Liapis, G. N. Yannakakis, M. J. Nelson, M. Preuss, and R. Bidarra, "Orchestrating game generation," *IEEE Transactions on Games*, vol. 11, no. 1, pp. 48–68, 2018.
- [2] G. N. Yannakakis and J. Togelius, *Artificial intelligence and games*. Springer, 2018, vol. 2.
- [3] N. Shaker, J. Togelius, and M. J. Nelson, *Procedural content generation in games*. Springer, 2016.
- [4] J. Togelius, A. J. Champanand, P. L. Lanzi, M. Mateas, A. Paiva, M. Preuss, and K. O. Stanley, "Procedural content generation: Goals, challenges and actionable steps." Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2013.
- [5] A. Summerville, S. Snodgrass, M. Guzdial, C. Holmgård, A. K. Hoover, A. Isaksen, A. Nealen, and J. Togelius, "Procedural content generation via machine learning (PCGML)," *IEEE Transactions on Games*, vol. 10, no. 3, pp. 257–270, 2018.
- [6] J. Liu, S. Snodgrass, A. Khalifa, S. Risi, G. N. Yannakakis, and J. Togelius, "Deep learning for procedural content generation," *Neural Computing and Applications*, vol. 33, no. 1, pp. 19–37, 2021.
- [7] D. Plans and D. Morelli, "Experience-driven procedural music generation for games," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 4, no. 3, pp. 192–198, 2012.
- [8] A. Naushad and T. Muhammad, "Condition driven adaptive music generation for computer games," *International Journal of Computer Applications*, vol. 64, no. 8, pp. 6–10, 2013.
- [9] S. Engels, T. Tong, and F. Chan, "Automatic real-time music generation for games," in *Proceedings of Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*, 2015.
- [10] A. K. Hoover, W. Cachia, A. Liapis, and G. N. Yannakakis, "Audioinspace: Exploring the creative fusion of generative audio, visuals and gameplay," in *International Conference on Evolutionary and Biologically Inspired Music and Art*. Springer, 2015, pp. 101–112.

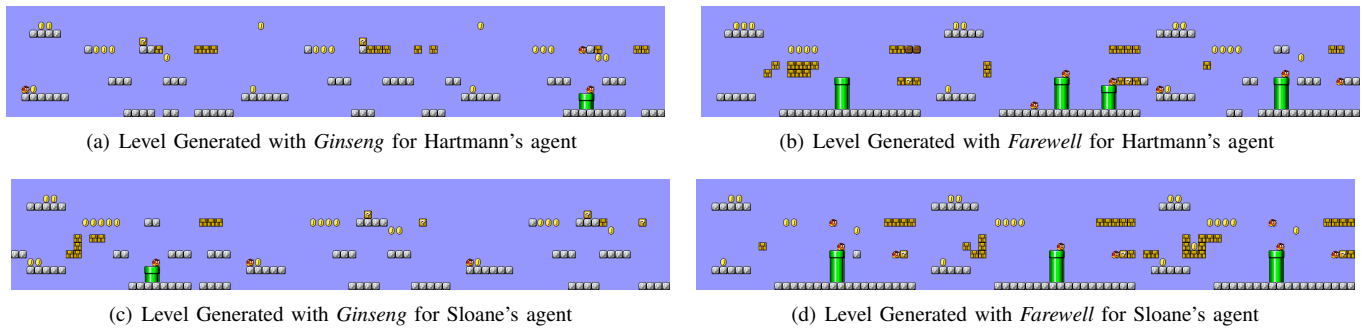


Fig. 7: Levels generated from an identical initial segment using the designer trained with  $F+C+P$  as reward, different musics and different agents as player. Levels are captured from the 21th segments, as the controller’s archive is fulfilled after 20 segments were generated.

[11] T. Shu, J. Liu, and G. N. Yannakakis, “Experience-driven PCG via reinforcement learning: A Super Mario Bros study,” in *Proceedings of IEEE Conference on Games*. IEEE, 2021, pp. 1–9.

[12] S. Earle, M. Edwards, A. Khalifa, P. Bontrager, and J. Togelius, “Learning controllable content generators,” in *Proceedings of IEEE Conference on Games*. IEEE, 2021, pp. 1–9.

[13] A. Khalifa, P. Bontrager, S. Earle, and J. Togelius, “PCGRL: Procedural content generation via reinforcement learning,” in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol. 16, no. 1. AAAI, 2020, pp. 95–101.

[14] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Proceedings of Advances in Neural Information Processing Systems*, vol. 27. Curran Associates, Inc., 2014.

[15] V. Volz, J. Schrum, J. Liu, S. M. Lucas, A. Smith, and S. Risi, “Evolving Mario levels in the latent space of a deep convolutional generative adversarial network,” in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2018, pp. 221–228.

[16] T. Shu, Z. Wang, J. Liu, and X. Yao, “A novel CNet-assisted evolutionary level repairer and its applications to super mario bros,” in *Proceedings of IEEE Congress on Evolutionary Computation*. IEEE, 2020, pp. 1–10.

[17] C. Donahue, Z. C. Lipton, and J. McAuley, “Dance dance convolution,” in *Proceedings of International Conference on Machine Learning*. PMLR, 2017, pp. 1039–1048.

[18] E. Halina and M. Guzdial, “TaikoNation: Patterning-focused chart generation for rhythm action games,” in *Proceedings of the Sixteenth International Conference on the Foundations of Digital Games*, 2021, pp. 1–10.

[19] A. Jordan, D. Scheffelowitzsch, J. Lahni, J. Hartwecker, M. Kuchem, M. Walter-Huber, N. Vortmeier, T. Delbrügger, Ü. Güler, I. Vatulkin *et al.*, “BeatTheBeat music-based procedural content generation in a mobile game,” in *Proceedings of IEEE Conference on Computational Intelligence and Games*. IEEE, 2012, pp. 320–327.

[20] D. Karavolos, A. Bouwer, and R. Bidarra, “Mixed-initiative design of game levels: Integrating mission and space into level generation,” in *Proceedings of Tenth International Conference on the Foundations of Digital Games*, 2015.

[21] P. W. Atmaja, R. Parlika, and F. Muttaqin, “Generating two-dimensional platformer game levels from storylines,” in *Proceedings of International Conference on Science and Technology*. Atlantis Press, 2018, pp. 1050–1056.

[22] R. Hunicke, “The case for dynamic difficulty adjustment in games,” in *Proceedings of the ACM SIGCHI International Conference on Advances in Computer Entertainment Technology*. ACM, 2005, pp. 429–433.

[23] P. Shi and K. Chen, “Learning constructive primitives for real-time dynamic difficulty adjustment in Super Mario Bros,” *IEEE Transactions on Games*, vol. 10, no. 2, pp. 155–169, 2017.

[24] D. Stammer, T. Günther, and M. Preuss, “Player-adaptive Spelunky level generation,” in *Proceedings of IEEE Conference on Computational Intelligence and Games*. IEEE, 2015, pp. 130–137.

[25] G. N. Yannakakis and J. Togelius, “Experience-driven procedural content generation,” *IEEE Transactions on Affective Computing*, vol. 2, no. 3, pp. 147–161, 2011.

[26] N. Shaker, G. N. Yannakakis, J. Togelius, M. Nicolau, and M. O’neill, “Evolving personalized content for Super Mario Bros using grammatical evolution,” in *Proceedings of Eighth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2012.

[27] P. M. Blom, S. Bakkes, C. T. Tan, S. Whiteson, D. Roijers, R. Valentini, and T. Gevers, “Towards personalised gaming via facial expression recognition,” in *Proceedings of Tenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2014, pp. 30–36.

[28] J. Togelius, S. Karakovskiy, and R. Baumgarten, “The 2009 mario ai competition,” in *IEEE Congress on Evolutionary Computation*. IEEE, 2010, pp. 1–8.

[29] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *Proceedings of International Conference on Machine Learning*. PMLR, 2018, pp. 1861–1870.

[30] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, and L. Sergey, “Soft actor-critic algorithms and applications,” *arXiv preprint arXiv:1812.05905*, 2018.

[31] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “OpenAI Gym,” *arXiv preprint arXiv:1606.01540*, 2016.

[32] A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” in *Proceedings of the International Conference on Learning Representations*, 2015.

[33] A. J. Summerville, S. Snodgrass, M. Mateas, and S. Ontanon, “The VGLC: The video game level corpus,” *arXiv preprint arXiv:1606.07487*, 2016.

[34] B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto, “librosa: Audio and music signal analysis in python,” in *Proceedings of the Fourteenth Python in Science Conference*, vol. 8. Citeseer, 2015, pp. 18–25.