

Level Generation for Angry Birds with Sequential VAE and Latent Variable Evolution

Takumi Tanabe

University of Tsukuba & RIKEN AIP
tanabe@bbo.cs.tsukuba.ac.jp

Jun Sakuma

University of Tsukuba & RIKEN AIP
jun@cs.tsukuba.ac.jp

Kazuto Fukuchi

University of Tsukuba & RIKEN AIP
fukuchi@cs.tsukuba.ac.jp

Youhei Akimoto

University of Tsukuba & RIKEN AIP
akimoto@cs.tsukuba.ac.jp

ABSTRACT

Video game level generation based on machine learning (ML), in particular, deep generative models, has attracted attention as a technique to automate level generation. However, applications of existing ML-based level generations are mostly limited to tile-based level representation. When ML techniques are applied to game domains with non-tile-based level representation, such as *Angry Birds*, where objects in a level are specified by real-valued parameters, ML often fails to generate playable levels. In this study, we develop a deep-generative-model-based level generation for the game domain of *Angry Birds*. To overcome these drawbacks, we propose a sequential encoding of a level and process it as text data, whereas existing approaches employ a tile-based encoding and process it as an image. Experiments show that the proposed level generator drastically improves the stability and diversity of generated levels compared with existing approaches. We apply latent variable evolution with the proposed generator to control the feature of a generated level computed through an AI agent's play, while keeping the level stable and natural.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks; Learning latent representations;**

KEYWORDS

Procedural Content Generation, AngryBirds, Sequential VAE, Latent Variable Evolution

ACM Reference Format:

Takumi Tanabe, Kazuto Fukuchi, Jun Sakuma, and Youhei Akimoto. 2021. Level Generation for Angry Birds with Sequential VAE and Latent Variable Evolution. In *2021 Genetic and Evolutionary Computation Conference (GECCO '21)*, July 10–14, 2021, Lille, France. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3449639.3459290>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '21, July 10–14, 2021, Lille, France

© 2021 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-8350-9/21/07...\$15.00

<https://doi.org/10.1145/3449639.3459290>

1 INTRODUCTION

Procedural content generation (PCG) automatically generates the content of games using an algorithm. The design of game content requires a large amount of knowledge and experience regarding the game. Therefore, manually creating new game content is an expensive task. Thus, PCG is a promising approach to minimize costs and generate new game content with different features, such as the difficulty of the game. PCG has been recently applied to various games for different purposes [14]. This study focuses on the automatic generation of game levels as a sub-category of PCGs.

Machine-learning-based PCG (PCGML) is a recent approach that has gathered research attention [18, 24]. This approach generates video game content using machine learning (ML) models trained on existing game content as the dataset. Human-designed game content satisfy implicit constraints such as whether a level is playable and has aesthetics. ML models learn such implicit constraints from a dataset without implementing them manually. Therefore, if the game content for a game is made available to the public, anyone can use it to build a game content generator without knowledge and experience regarding the game.

PCGML includes approaches that utilize n-gram [4] with a graphical probability model [9], generative adversarial networks (GAN) [7, 26, 27], and variational autoencoders (VAEs) [25]. Volz et al. [27] applied the Wasserstein GAN (WGAN) to generate levels for *Super Mario Bros* (SMB). Torrado et al. [26] successfully generated playable levels with a small dataset using conditional embedding self-attention GAN and the bootstrapping method. Thakkar et al. [25] compared the autoencoder and VAE on game-level generation and demonstrated that the VAE can generate more complex levels.

An advantage of PCGML using a deep generative models (DGMs), such as GAN or VAE, is the possibility of controlling the characteristics of the generated game contents by applying latent variable evolution (LVE) [7, 27]. In these approaches, the game content generator is modeled as a map from a real-valued latent vector to a game content, and is trained on an existing dataset. The latent vector is optimized such that the generated game contents have the desired features, which can be evaluated through a simulation using AI or human game players. For example, Volz et al. [27] first trained a GAN to generate game levels for SMB and then applied LVE to control the difficulty of the generated levels, such as the number of jumps that an AI player can make. With DGM and LVE, one can generate game contents that satisfy the constraints imposed in a training dataset implicitly by human designers, while

realizing some implicit features that may not exist in the training dataset.

Although there are strict constraints on the game level in that it must be playable, unlike the general targets of DGMs (e.g., audio and images), DGMs have been utilized in loosely constrained game (e.g., SMB) that are playable even if the blocks are placed somewhat haphazardly. Training high-quality DGMs requires a large amount of data, whereas most games can only train low-quality models that produce locally incomplete levels due to the limited data available. Therefore, it is difficult for a low-quality generative models to generate levels that satisfy the constraint of being playable [24]. A drawback of previous studies is that they only focused on games with loosely constrained levels, even though many game levels have strict constraints.

In this study, we consider PCGML-based level generation for *Angry Birds*, which is a casual action puzzle game. *Angry Birds* is often targeted in PCG as a challenging game domain for level generation¹. The game applied in the existing study was a game with tile-based levels, whereas the *Angry Birds* are described by an XML file containing various object types with different sizes, measured coordinates, and rotations. Therefore, the design patterns of levels are massive compared to tile-based levels, and it is difficult to generate stable levels. The presence of gravity in the level space, combined with the previous difficulty, makes it challenging to generate playable level. This is because a slight shift in the position of the blocks can cause the level to collapse immediately after the start of the game due to the force of gravity, making the generated level unplayable. Earlier studies have focused on tile-based game levels with no gravity constraints; therefore, it is difficult to apply existing methods to this game domain, as described in Section 2. To the best of our knowledge, there are no examples of level generation methods where DGMs have been applied to *Angry Birds*, or to games with similar playability constraints.

To deal with the aforementioned difficulties, we propose building a level from the bottom up by gradually dropping objects considering the effect of gravity. We encode a game level as sequence data, where the time step indicates the order in which the objects are dropped. The types, horizontal positions, and rotations of the dropped objects are also specified in each time step. Each combination of this specification is considered as a word, and one sentence corresponds to a level. In sequence encoding, the vertical positions of objects are automatically determined by the order and height of dropping objects, which drastically improve the stability of levels under the gravity. A long short-term memory (LSTM)-based VAE is used as the generative model. Furthermore, word embedding techniques are employed to improve the stability of the generated levels. Then, LVE is applied to design a stable level with the desired features. Our implementation is available publicly².

Three experiments are conducted to demonstrate the usefulness of the proposed approach. First, we compare the stability of the levels generated by the proposed method with those of various models such as tile-based (image-based) generation models and various encoding methods, demonstrating that our encoding method is effective in generating stable levels with high probability.

Second, we compare the diversity of the levels generated by the existing level generator that is not PCGML-based with those generated by the proposed generator trained on the levels generated by the existing one. When the diversity of the levels generated is extremely low, it is highly likely that stable levels can be generated. In this experiment, we have shown that the levels generated by the proposed method can generate a variety of levels while maintaining high stability. Third, to demonstrate the advantages of using our generative model-based approach, we apply LVE to the trained model. We successfully optimized the number of the specific types of objects in the level and the objective function computed through the AI agent's play.

2 ANGRY BIRDS

Our game domain is a clone of *Angry Birds*, called *Science Birds* [5] for the visualization and simulation of levels. It is a very active game domain in the field of procedural level generation, where level generation competitions are held every year [23].

In this game, the player destroys all the pigs present in a level by shooting a limited number of birds with a slingshot. The game level contains different types of objects with different shapes and sizes. The objects in this game can be roughly classified into platform, TNT blocks, pig blocks, and regular blocks.

The game level of *Angry Birds* is specified by a list of objects. Each object contains the type (e.g., small square stone), vertical and horizontal coordinates, and angle. There is gravity in the level, and most types of blocks in the level are affected by gravity. Hence, blocks must be stacked stably under the gravity, by considering the differences of blocks in terms of size and shape.

There are two main difficulties in developing a ML-based level generation method for *Angry Birds*. (i) It has a very high degree of freedom in the level design. The existing level generation techniques that employ DGMs have been applied to tile-based game domains wherein the levels are clearly divided into meshes and all blocks have the same size or the same unit size, and the blocks can be set by specifying the index of the level array. Therefore, the existing approaches such as [25–27] encode the tile-based level as an image and successfully use DGMs for image processing. However, the levels of *Angry Birds* are specified by real numbers, and a large number of meshes are needed to represent the exact stage. The different block sizes can easily result in overlapping blocks, which may lead to an unplayable level. (ii) Another difficulty is that the conditions for reaching a playable level are severe because of the existence of gravity. A small error can cause a collapse of the entire level immediately after the game starts, as shown in Figure 1.

Existing level generation approaches for *Angry Birds* are designed using domain knowledge, and the search-based approaches are employed. Ferreira et al. [6] used a genetic algorithm to find combinations of basic blocks and pre-defined baseline structures that constitute stable levels. Stephenson et al. [21] predefined an algorithm to generate building blocks that are likely to be stable. Changes in pig placement, location of additional blocks, and block materials are updated using heuristics to increase level robustness, and changes in block types to increase level diversity. This was extended in [22], which optimized the parameters of the level generator to maximize the utility computed by the simulation using AI

¹<https://aibirds.org/level-generation-competition.html>

²<https://github.com/yoshinobc/Level-Generation-for-Angry-Birds-with-Sequential-VAE-and-Latent-Variable-Evolution>




Figure 1: Examples of *Angry Birds* levels. Left: A stable level. Right: Unstable level, where a small stone object is removed from the left figure. This corrupts after the game starts.

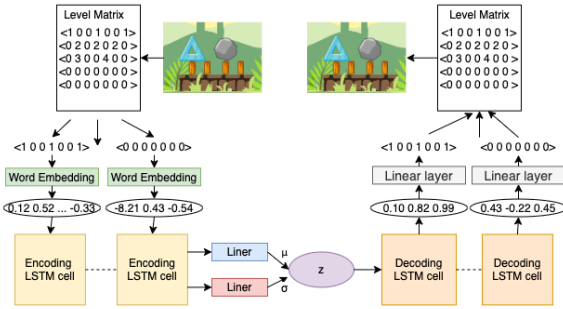


Figure 2: Flowchart of the proposed approach

agents. Abdullah et al. [1] proposed an algorithm to create levels that can be destroyed in one shot by combining human-designed baseline structures. These methods rely on domain knowledge to build stable baseline structures with different objects under the gravity.

3 PROPOSED APPROACH


We aimed at developing a DGM-based level generation for game domains where the difficulty in generating stable levels described in Section 2 exists. We focus on *Angry Birds* as an example of such game domains in this study.

We propose a DGM-based level generation for *Angry Birds*. The objective is to replace the engineering process of the existing level generation approaches by training a generative model from a dataset and enabling LVE on the top of the trained generator. To utilize an efficient algorithm designed for a general-purpose black-box continuous optimization in the LVE stage, for example, CMA-ES [11–13], we employ a generative model that uses a continuous latent vector as an input.

Figure 2 shows the overall diagram of the proposed VAE. Each level is encoded as sequence data (Section 3.1). A sequence is treated as a sentence, and each entry of the sentence is treated as a word (Section 3.2). Word embedding is applied (Section 3.3). Our generative model is a sequential VAE [2], which is trained to resemble existing levels in a dataset (Section 3.4). The LVE explores the latent space of the trained VAE by optimization, for a level with the desired features (Section 3.5). The technical details of the level encoding and decoding are presented in Appendices A and B.



(a) Dense encoding



(b) Sparse encoding

Figure 3: Image-based encoding approaches



Figure 4: Proposed encoding method

3.1 Sequential Encoding

We propose encoding a level as a sequence data, $S = (s_1, \dots, s_{\text{MAX_ROW}})$. We refer to this as a *level matrix*. Unlike the tile-based encoding displayed in Figure 3 where a level is divided into a predefined mesh and objects are placed on cells, our encoding strategy models the level from the bottom to the top. The time step, t , indicates the order of dropping objects, and the type and horizontal positions of the dropping objects are specified by s_t . Thus, we incorporate the role of gravity naturally. All objects are forcefully placed on the ground or on other objects. The maximal time step is set to the maximal length in the training data, which is $\text{MAX_ROW} = 30$ in our experiments.

Let \mathcal{O} denote the set of possible object types. The type of an object is determined by the combination of a type of object (e.g., wood) and its predefined rotation angle. We have $|\mathcal{O}| = \text{N_TYPES} = 61$ types in total. The horizontal coordinate of a level is divided into a uniform mesh of size $\text{MAX_COL} = 94$ in advance. In other words, $s_t = (o_{t,1}, \dots, o_{t,\text{MAX_COL}})$ for each t , and $o_{t,i} \in \mathcal{O}$ is an object dropped at the i -th horizontal position at time step t . Because objects in \mathcal{O} are different in size, the vertical coordinates of $o_{t,i}$ and $o_{t,j}$ are not necessarily the same. These coordinates are determined by physically dropping the specified object at the level containing the objects placed in the previous time steps.

Figure 4 depicts this encoding. We let \mathcal{D}_S denote the training dataset of the levels that are processed by the proposed encoding.

3.2 Word-based Representation

We propose to further process a level matrix, S , as if it is a natural language sentence. A level matrix is regarded as a sentence, each row, s_t , of a level matrix is regarded as a word, and each column, $o_{t,i}$ in s_t , is regarded as a letter. The number of alphabets is N_TYPES ; a word is composed of MAX_COL letters. A level matrix, S , is mapped to a sequence $\mathbf{W} = (\mathbf{w}^{(t)})_{t=1}^T$ of one-hot vectors $\mathbf{w}^{(t)}$ representing a word in the vocabulary. Vocabulary $\mathcal{W} = \{s : s \in S \text{ for } S \in \mathcal{D}_S\}$ is a set of all rows appearing in the level matrices in the training dataset. Indexing is performed, and the one-hot vector of a word represents the word index. The

number of words in a sentence is denoted by T , and $T = \text{MAX_ROW}$ in our case. Let \mathcal{D} denote the dataset of the word-based sequence representation of the level matrices in \mathcal{D}_S .

The rationale behind the choice of word-based representation is as follows. In text generation using deep neural networks, *word-based* modeling is more likely to lead to better performance than a *character-based* modeling [17]. In character-based modeling, a neural network must learn combinations of characters that exist in the vocabulary, and correlations between words in a sentence. In our case, a word that does not exist in the vocabulary corresponds to a pattern of blocks that does not appear in a training dataset. On the one hand, it may be helpful to have a large variation in the generated levels after training. On the other hand, patterns that do not appear in a training dataset are likely to be invalid (objects that are located too closely can overlap) or unstable. Based on the above considerations, we employ a word-based approach.

3.3 Word Embedding

Word embedding is applied to obtain a real vector representation, \mathbf{x} , of a word \mathbf{w} . A difficulty in training a VAE for sequence data is its high dimensionality. A naive representation of a word is a one-hot representation, where the length of the vector is $|\mathcal{W}|$. It is sparse and high dimensional, and it is difficult to train the VAE adequately for such data. To easily train the VAE, we utilize word embedding that provides dense and low-dimensional representation.

We employ word2vec with the continuous Bag-of-Words model [19]. The model is composed of two linear layers, \mathcal{E} and \mathcal{F} , where $\mathcal{E} : \mathbf{w} \mapsto \mathbf{x}$ maps a word to its corresponding embedding, $\mathbf{x} \in \mathbb{R}^{\text{DIM}_X}$ and $\mathcal{F} : \mathbb{R}^{\text{DIM}_X} \rightarrow \mathbb{R}^{|\mathcal{W}|}$. The conditional probability of $\mathbf{w}^{(t)}$ given $(\mathbf{w}^{(t-m)}, \dots, \mathbf{w}^{(t-1)}, \mathbf{w}^{(t+1)}, \dots, \mathbf{w}^{(t+m)})$ is modeled by using \mathcal{F} , where the probability vector is given by the SOFT-MAX of $\mathcal{F}(\sum_{i=1}^m \mathcal{E}(\mathbf{w}^{(t-i)}) + \mathcal{E}(\mathbf{w}^{(t+i)}))$. The parameters of \mathcal{E} and \mathcal{F} are trained to maximize the conditional probability on a training dataset. Then, $\mathbf{x}^{(t)} = \mathcal{E}(\mathbf{w}^{(t)})$ is used as the embedding of the word $\mathbf{w}^{(t)}$. The sequence of the embedding vectors is denoted by $\mathbf{X} = (\mathbf{x}^{(t)})_{t=1}^T$. For simplicity, we consider $\mathbf{X} = \mathcal{E}(\mathbf{W})$ in the following.

3.4 Sequential Variational Autoencoder

We model the generation of a sequence \mathbf{W} that represents a game level as a variant of the VAE, i.e., sequential VAE [2]. A VAE consists of two models: an encoder and a decoder. The encoder is trained to map a high-dimensional feature vector to a distribution of a low-dimensional latent vector, and the decoder is trained to recover an original feature vector from a latent vector. Both the encoder and decoder networks are modeled by LSTM networks, which can handle sequential data, because our data are sequential data. The encoder network is a map from a sequence $\mathbf{X} = (\mathbf{x}^{(t)})_{t=1}^T$ to the mean vector and the log of the coordinate-wise variance of a Gaussian distribution, which models the conditional probability, $q_\phi(\mathbf{z} | \mathbf{X})$, of a latent vector, $\mathbf{z} \in \mathbb{R}^{\text{DIM}_Z}$, associated with the given sequence, \mathbf{X} . The decoder network models the conditional probability, $p_\theta(\mathbf{W} | \mathbf{z})$ of a sequence $\mathbf{W} = (\mathbf{w}^{(t)})_{t=1}^T$ for one-hot word representations, given the latent representation, \mathbf{z} .

The sequential VAE is trained by minimizing the loss function defined by β -VAE [15]. The loss function consists of two components: reconstruction loss and Kullback–Leibler (KL) loss. The reconstruction loss is defined as

$$\mathcal{L}_{\text{rec}}(\mathcal{D}) = \frac{1}{|\mathcal{D}|} \sum_{\mathbf{W} \in \mathcal{D}} \ln p_\theta(\mathbf{W} | \mathbf{z}_{\mathbf{W}}), \quad \mathbf{z}_{\mathbf{W}} \sim q_\phi(\cdot | \mathcal{E}(\mathbf{W})).$$

This measures the likelihood of the input sequence to the encoder being reconstructed by the decoder. The KL loss is defined as

$$\mathcal{L}_{\text{KL}}(\mathcal{D}) = \frac{1}{|\mathcal{D}|} \sum_{\mathbf{W} \in \mathcal{D}} D_{\text{KL}}(q_\phi(\mathbf{z} | \mathcal{E}(\mathbf{W})) \| \mathcal{N}(\mathbf{0}, \mathbf{I})),$$

where D_{KL} denotes the KL divergence. This measures the distance from the distribution of latent vectors to the Gaussian prior. Finally, the loss is defined as the weighted sum, $-\mathcal{L}_{\text{rec}}(\mathcal{D}) + \beta \cdot \mathcal{L}_{\text{KL}}(\mathcal{D})$. Based on [2], we adopt two techniques—KL annealing and word drop—when training the model. These techniques have been introduced to mitigate the issue in which the trained decoder ignores the latent vector. For details, we refer to [2, 15].

3.5 Level Generation

We define our level generator, $\mathcal{G} : \mathbf{z} \mapsto \mathbf{W}$, as a deterministic map. Given a latent vector, \mathbf{z} , $\mathcal{G}(\mathbf{z})$ is defined by the maximizer, \mathbf{W} , of the trained decoder, $p_\theta(\mathbf{W} | \mathbf{z})$. The levels generated by \mathcal{G} are expected to satisfy *implicit constraints imposed on the training dataset*. Here, the generated levels are expected to be stable and playable because the training data are stable and playable.

To explore levels that have a desired features, we perform LVE. Let $f(\mathbf{W})$ be an objective function. The objective function value of a level \mathbf{W} can be computed directly from \mathbf{W} , using the results of the AI agents' play or evaluated by human players on a game renderer. We pose the following optimization problem:

$$\min_{\mathbf{z}} f(\mathcal{G}(\mathbf{z})) \quad \text{subject to } \mathbf{z} \in \mathbb{Z}, \quad (1)$$

where $\mathbb{Z} \subseteq \mathbb{R}^{\text{DIM}_Z}$ denotes the search space. Because the generator is trained only for $\mathbf{z} \in \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\mathcal{G}(\mathbf{z})$ may not be trained for \mathbf{z} , which is unlikely for $\mathcal{N}(\mathbf{0}, \mathbf{I})$. Considering this point and the ease of handling constraints, we set $\mathbb{Z} = [-3, 3]^{\text{DIM}_Z}$.

4 EXPERIMENTS

We compared the probability of producing a playable levels with generators using different encoding methods. The comparison results indicated that word-based encoding methods are more suitable than char-based and image-based encoding methods. Further, we investigated the diversity of levels produced by different generators. Based on the results of this comparison, we demonstrated that our proposed encoding method generates more diverse levels than other encoding methods.

4.1 Common setting

Dataset. We used a dataset consisting of 200 levels, where 180 levels were considered as the training dataset and 20 levels as the validation dataset. The validation set was used to adjust the training hyper-parameters. The dataset size was significantly smaller than those typically used for DGM training; however, it was a realistic number in level generation because datasets used in level

Table 1: Numbers of unique uni-gram and bi-gram types present in the levels generated by different methods.

	Training	Proposed	W/o WE	Char(D)	Char(S)
Uni-gram	1503	1376	233	88	426
Bi-gram	1832	2737	267	93	848

Table 2: Number of stable levels among 100 levels generated from each generator. We consider a level to be stably generated if it does not fall after the game starts.

Proposed	W/o WE	Char(D)	Char(S)	Image(D)	Image(S)
96	93	0	10	6	0

generation are created manually.³ The dataset was generated from IratusAves [21]⁴, which was the generator that won the CIG 2018 Level Generation Competition. The vocabulary size, $|W|$, in the training dataset was 1503. Although it may not be very practical, we generated a larger dataset of 10000 levels and trained our model on it for testing; the results are provided in Appendix C.

Model. The output of the word embedding had a dimension $\text{DIM}_X = 50$. The dimension of the latent vector in the VAE was $\text{DIM}_Z = 60$. The number of hidden neurons for the encoder and decoder LSTM networks was 400. The VAE was trained on the training data with 500 epochs. The word drop rate was set to 0.3, and the KL loss was ignored in 250 epochs.

Baselines. We compared the following approaches:

- Proposed: as described in Section 3
- W/o WE: Word-based encoding without word embedding
- Char(D): Character-based dense encoding
- Char(S): Character-based sparse encoding
- Image(D): Image-based dense encoding
- Image(S): Image-based sparse encoding

In dense encoding (Figure 3a), cells corresponding to an object are filled with the object index, whereas in sparse encoding, the object index is located only at the center cell of the object (Figure 3b). In the image-based encoding, a level is divided into a uniform mesh, and each cell has an object type index. We employed WGAN-GP [8] for the image-based approaches, that were used in the existing level generator [27]. In WGAN-GP, the row and column sizes of a level were set as 96. There was no difference between dense and sparse encoding in word-based encoding because they were transformed into word indices. For image-based encoding approaches, an imperfect level matrix (overlapping objects and missing and extra cells for an object) can be repaired to some extent by manually designed procedure.

³In Angry Birds, the size of the level corpus created by Zafar et al. [28] was 200. In other games, the size of the dataset used to generate levels for *Lode Runner* was 150 [25], and that for *SMB* was 173 [27].

⁴We used the code provided at <https://github.com/stepmat/IratusAves>.

4.2 Diversity

We compared the diversity of the level matrices computed from the training dataset and those generated by the proposed method, W/o WE, Char(D), and Char(S).

A level matrix was generated by sampling a latent vector from the multivariate normal distribution, $\mathcal{N}(\mathbf{0}, \mathbf{I})$, and feeding it to the trained decoder. We produced a 1000 level matrix using each method.

We measured the diversity of level matrices using n-gram, which is used in natural language processing. An n-gram is a contiguous sequence of N words or characters. Since our generator is word-based, we use word as the unit. The diversity measure for n-grams is called distinct-N (a diversity-promoting objective function for neural conversation models). We counted the number of unique unigrams ($N = 1$) and bigrams ($N = 2$). The higher the number of unigrams present in the generated levels, the richer are the word (a group of blocks dropped at once in a single trial) patterns are in the generated levels. The more bigrams there are in the generated levels, the higher is the number of levels that can be generated with various combinations of words.

Table 1 summarizes the numbers of unique unigrams and unique bi-grams included in the training dataset and 1000 level matrices generated by the trained generators. Word-based encoding cannot generate unigrams that do not appear in the training dataset; however, it is possible to generate bi-grams that do not appear in the training dataset.

The results indicated that approximately 92% of uni-gram types can be reproduced, and that the variation of bigrams was as high as 149% of the training dataset. These values, especially those for bigrams, may even increase as the number of generated levels increases. For W/o WE, char(S), char(D), the diversity of the generated levels was very low for both unigrams and bigrams, probably because of the difficulty in learning.

4.3 Stability

We evaluated the stability of the levels generated by the proposed level generator and those generated using different encoding approaches. We verified our expectations: 1) the proposed sequential encoding approach is advantageous in generating stable levels over image-based encoding approaches and 2) the word-based encoding approaches produce more stable levels than character-based encoding approaches. Stability is measured by the probability of successfully generating a level that does not collapse without any action.

Table 2 lists the number of stable levels among the 100 levels generated by each approach. Sequential encoding drastically improved the stability of the generated levels over image-based approaches. The proposed approach succeeded in generating stable levels with a probability of about 96%. Further, W/o WE generated stable levels with a high probability of 93%. The levels produced by the generators of character-based and image-based encoding methods were not stable.

Figure 5 shows example levels generated by these approaches. The levels generated by image-based approaches were unstable. Character-based approaches often failed to stack objects (objects moved left or right as they were stacked on each other). Here, we

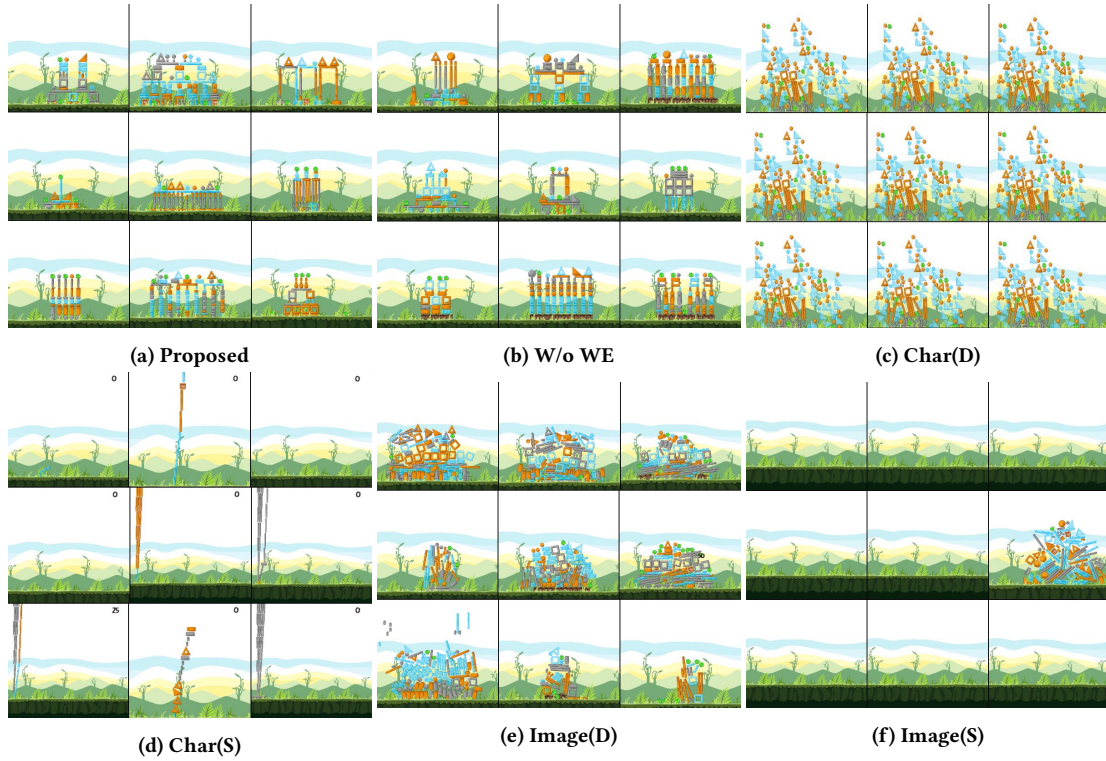


Figure 5: Levels generated by different generators

observed the difficulty in training the correlation between distant characters. The levels were more likely to be stable for word-based approaches, especially for the proposed approach.

4.4 Discussion

The results of the aforementioned experiments indicated that our proposed generator and W/o WE—both word-based encoding strategies—generated more stable levels than char-based (Char(D) and Char(S)) and image-based (Image(S) and Image(D)). Further, the diversity of the levels generated by the proposed approach was higher than that generated by other sequence-based encoding strategies (W/o WE, Char(D), and Char(S)).

In the experiment described in Section 4.2, character-based embedding can theoretically generate words that do not appear in the training dataset; hence, the number of unique unigrams of Char(D) may be larger than that of the training dataset. However, in this experiment, the number of unigrams and bigrams was smaller than that in the training dataset. This is probably because the dataset size was relatively small compared to those typically used to train DGMs, which resulted in an unsuccessful learning procedure. Appendix C details the results on the larger dataset.

Word embedding had the effect of dimensional compression. Because the dataset used in this study was small (the number of dimensions is 1503), it was not too large even without word embedding. Therefore, in this experiment, there was no considerable difference between the model without word embedding and the generator with word embedding. However, on the large dataset

outlined in Appendix C, the level of stability was lower without word embedding. For large datasets, the loss during the training of the sequential VAE for the proposed method was approximately 10^3 times smaller than that for the W/o WE. This implied that the word embedding makes model training more tractable.

In the experiment described in Section 4.3, generators of character-based and image-based encodings unlike those of word-based encodings could rarely produce stable levels. Word-based encoding was different from char-based encoding in that the combinations of blocks were determined in advance as words. Because word-based encoding determines only the order of combination, it can generate stable levels. However, char-based encoding determines the combination of blocks without using the word. Therefore, compared to word-based encoding, char-based encoding generated fewer stable levels. Unlike the image-based encoding method, the word-based encoding method generated levels by dropping blocks from the top. This solved the problems of overlapping blocks and blocks floating apart, thereby achieving stability for word-based encoding compared to that for image-based encoding.

5 LATENT VARIABLE EVOLUTION

An advantage of using DGMs for level generation is that LVE [7, 27] can be employed after training the generator to identify levels with certain characteristics. However, existing studies on sequential VAE have reported that the latent vector tends to be ignored, and the variation of outputs is created by the randomness




Figure 6: Single run results of LVE without simulation. Average (solid line) and standard deviation (band) of N_{Pigs} (left) and N_{TNTs} (right) over ten levels generated using the best solution for each generation.

of \mathcal{G} itself [20]. Thus, \mathcal{G} is not suitable for LVE, because we cannot control the features of the generated levels. Although we adopted techniques to mitigate this issue in training the sequential VAE, the applicability of LVE with the proposed level generator must be evaluated. In this section, we demonstrate the applicability by performing LVE with four different objectives.

5.1 LVE without Simulation

The first experiment was to maximize the number of specific blocks, namely, i.e., the number N_{Pigs} of pig blocks and the number N_{TNTs} of TNT blocks. These are the most important blocks in *Angry Birds* and determine the the game's interest.

In this experiment, instead of directly optimizing the latent vector, z , by solving (1), we parameterized the latent vector with $\alpha \in \mathbb{R}^{\text{DIM}_Z}$ and $\beta \in \mathbb{R}^{\text{DIM}_Z}$ as

$$z \sim \mathcal{N}(\beta, \alpha \cdot I) .$$

After optimization, we utilized the solution, (α^*, β^*) , to generate multiple latent vectors, z_1, \dots, z_m , which were considered as input for the generator, thus leading to a variety of levels, $\mathcal{G}(z_1), \dots, \mathcal{G}(z_m)$. Therefore, this formulation allowed variation in levels after the LVE process. We minimized the following objective function:

$$F(\alpha, \beta) = \frac{1}{m} \sum_{i=1}^m f(\mathcal{G}(z_i)) \quad z_i \sim \mathcal{N}(\beta, \alpha \cdot I) ,$$

where f denotes the negative of N_{Pigs} or the negative of N_{TNTs} . We set $m = 30$ in this experiment.

CMA-ES was used to optimize α and β with a population size of 60. We set the search space for (α, β) to $[0, 2] \times [-3, 3]^{\text{DIM}_Z}$. The software library pycma [10] was used. We used the proposed generator, \mathcal{G} , trained in Section 4.

Figure 6 shows the results of the optimization. The average and standard deviation of N_{Pigs} and N_{TNTs} were computed at 10 levels, $\mathcal{G}(z_1), \dots, \mathcal{G}(z_{10})$, using independent and $\mathcal{N}(\beta_t^*, \alpha_t^* \cdot I)$ distributed random vectors, z_1, \dots, z_{10} , where (α_t^*, β_t^*) was the best solution generated at generation t . We observed that the numbers of these blocks were increased successfully. This implied that the trained generator does not ignore the latent vectors, and the features (here, N_{Pigs} and N_{TNTs}) can be controlled by the LVE.

Figure 7 shows the levels generated during the optimization. Here, two points need to be considered. First, although our objective was to maximize the number of pig blocks, the number of other




Figure 7: Results of the maximization of N_{Pigs} . Levels generated at the 20th (top left), 40th (bottom left), and 100th (right) generations are displayed. There are 5, 6, and 13 pig blocks in the levels, respectively.

blocks increased as well. This is because the training dataset did not include levels that contained pig blocks alone. This is promising because we would like to generate a level that can be naturally compared to the training dataset but has a features that are not necessarily included in the dataset. Second, the levels generated by the LVE process were more or less stable. However, there are levels produced during the LVE process that are not fully stable, such as the level generated at the 40th generation in Figure 7. Although the generator was trained to output stable levels, employing a constraint-handling technique to penalize unstable levels during the LVE process may be desired.

5.2 LVE with Simulation

Next, we demonstrated the LVE using a simulation with an AI agent⁵. We designed two objective functions that were computed through the AI agents play.

The first objective function aimed at measuring the difficulty of a level was computed as

$$\max(5 - N_{\text{Birds}}, N_{\text{RemBirds}}) \times 10 + N_{\text{Blocks}} , \quad (2)$$

where N_{Birds} , N_{Blocks} , and N_{RemBirds} represent the numbers of birds at the level at the start of the game, blocks at the level at the start of the game, and birds at the level after the game ends, respectively. The number of birds at the start of the game is automatically determined by the number of pigs and blocks present at the level, which is similar to that in *IratusAves*. We expected the optimized level to have $N_{\text{Birds}} \geq 5$, and they should all be used by the AI agent (N_{RemBirds}). The number of blocks in the level was minimized; however, the priority was placed on the first term (because of the coefficient of 10).

The second objective function aimed at measuring the aesthetics of a level by determining how well the level retained its original shape when all pig blocks were destroyed; the level was then

⁵<https://gitlab.com/aibirds/sciencebirdsframework>




Figure 8: Single run results of LVE with simulation. The average (solid line) and standard deviation (band) of the objective values, Equation (2) (left) and Equation (3) (right), over a population of 60 for each generation are displayed.

cleared, which was calculated by

$$\max(60 - N_{\text{Blocks}}, N_{\text{Blocks}} - N_{\text{RemBlocks}}) \times 10 - N_{\text{Pigs}} \quad (3)$$

where $N_{\text{RemBlocks}}$ represents the number of blocks at the level after the game ends.

CMA-ES with a boundary constraint implemented in *pycma* was used to optimize (1) with a population size of 60. We used the proposed generator, \mathcal{G} , trained in Section 4.

The results are presented in Figure 8. The average and standard deviation of the objective values over 60 candidate solutions generated at each iteration are displayed. Figure 9 shows the levels generated at the last generation (Equation (2) is 60 generations; Equation (3) is 100 generations). The objective functions were successfully minimized, and the generated levels were stable; although subjective, they looked natural. As shown in Figure 9a, the objective was reflected in the generated level. Equation (2) was designed to minimize the number of birds left after the AI agent plays. Further, the number of blocks was minimized, which was expected to result in a simple but challenging level. In fact, at this level, the pigs were grouped in four discrete positions, and the number of birds was also four, so it is probably necessary to use all the birds to successfully shoot all the pigs. For Figure 9b, the level was optimized to minimize the blocks that the birds have broken after the AI agent plays the game. In fact, the level after the agent has played the game showed that most of the blocks were unbroken, as shown in Figure 10.

6 CONCLUSION AND FUTURE WORK

We proposed a level generator for *Angry Birds*. To overcome the difficulties in generating a level that is stable under gravity, we encoded the level as sequence data and train a sequential VAE for level generation. The experimental results showed that the trained generator can generate unigrams as diverse as the levels of the training data, levels with more bigrams than the training data, and stable levels with high probability. We developed a level generator for *Angry Birds*, in which, it is difficult to generate playable levels using existing DGMs. Unlike existing non-ML-based level generators that target *Angry Birds*, we did not rely on domain knowledge. We expect that this approach can be generalized to other game domains with gravity at this level.

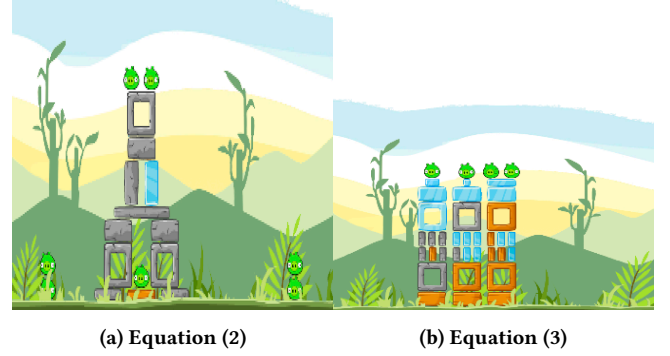


Figure 9: Results of minimization of Equation (2) (left) and Equation (3) (right). Levels generated at the 60th and 100th generation are displayed.




Figure 10: Level after the AI agent's play on Figure 9b.

Two possible directions to extend the proposed approach are as follows. One is designing a constraint-handling technique to penalize unstable levels during the LVE process. Although the proposed level generator is trained to output stable levels when the latent vector is generated from the normal distribution, the output can be unstable when performing LVE, as shown in Figure 7. The other direction is to develop a methodology to evaluate the level in a human-in-the-loop manner. Designing an objective function to reflect a qualitative measure such as aesthetics is a difficult task.

The PCG is not only useful for game content designers, but also for improving the robustness of AI agents based on reinforcement learning. In reinforcement learning, it has been reported that if the environment used for training is fixed, it will overfit the specific training environment and fail to generalize to the modified environment [16]. A promising approach to learn robust policies is to train agents at many different game levels [3]. The proposed level generator was used to train a robust *Angry Birds* agent. Moreover, we expect that LVE will allow us to locate the weakness of a learning agent and actively train it by generating levels at which the agent is poor. This is another direction for our future work.

ACKNOWLEDGMENTS

This work was partially supported by JSPS KAKENHI Grant Number 19H04179.

REFERENCES

- [1] Febri Abdullah, Pujana Paliyawan, Ruck Thawonmas, and Fitra A Bachtiar. 2020. Generating Angry Birds-Like Levels With Domino Effects Using Constrained Novelty Search. In *2020 IEEE Conference on Games (CoG)*. IEEE, 698–701.
- [2] Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew M. Dai, Rafal Jozefowicz, and Samy Bengio. 2016. Generating Sentences from a Continuous Space. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*. 10–21.
- [3] Karl Cobbe, Oleg Klimov, Chris Hesse, Taehoon Kim, and John Schulman. 2019. Quantifying Generalization in Reinforcement Learning. In *ICML 2019 : Thirty-sixth International Conference on Machine Learning*. 1282–1289.
- [4] Steve Dahlskog, Julian Togelius, and Mark J. Nelson. 2014. Linear levels through n-grams. In *Proceedings of the 18th International Academic MindTrek Conference on Media Business, Management, Content & Services*. 200–206.
- [5] Lucas Ferreira and Claudio Toledo. 2014. A Search-based Approach for Generating Angry Birds Levels. In *Proceedings of the 9th IEEE International Conference on Computational Intelligence in Games (CIG'14)*.
- [6] Lucas Ferreira and Claudio Toledo. 2014. A search-based approach for generating Angry Birds levels. In *2014 IEEE Conference on Computational Intelligence and Games*. 1–8.
- [7] Edoardo Giacomello, Pier Luca Lanzi, and Daniele Loiacono. 2019. Searching the Latent Space of a Generative Adversarial Network to Generate DOOM Levels. In *2019 IEEE Conference on Games (CoG)*. 1–8.
- [8] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. 2017. Improved training of wasserstein GANs. In *NIPS'17 Proceedings of the 31st International Conference on Neural Information Processing Systems*. 5769–5779.
- [9] Matthew Guzdial and Mark O Riedl. 2016. Game Level Generation from Gameplay Videos. In *AIIDE*. 44–50.
- [10] Nikolaus Hansen, Youhei Akimoto, and Petr Baudis. 2019. CMA-ES/pycma on Github. Zenodo, DOI:10.5281/zenodo.2559634. (Feb. 2019). <https://doi.org/10.5281/zenodo.2559634>
- [11] Nikolaus Hansen, Sibylle D. MÅeller, and Petros Koumoutsakos. 2003. Reducing the Time Complexity of the Derandomized Evolution Strategy with Covariance Matrix Adaptation (CMA-ES). *Evolutionary Computation* 11, 1 (2003), 1–18. <https://doi.org/10.1162/106365603321828970> arXiv:<https://doi.org/10.1162/106365603321828970>
- [12] N. Hansen and A. Ostermeier. 1996. Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. In *Proceedings of IEEE International Conference on Evolutionary Computation*. 312–317.
- [13] Nikolaus Hansen and Andreas Ostermeier. 2001. Completely Derandomized Self-Adaptation in Evolution Strategies. *Evol. Comput.* 9, 2 (June 2001), 159–195. <https://doi.org/10.1162/106365601750190398>
- [14] Mark Hendrikx, Sebastiaan Meijer, Joeri Van Der Velden, and Alexandru Iosup. 2013. Procedural content generation for games: A survey. *ACM Transactions on Multimedia Computing, Communications, and Applications* 9, 1 (2013), 1–22.
- [15] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. 2017. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *ICLR 2017 : International Conference on Learning Representations 2017*.
- [16] Niels Justesen, Ruben Rodriguez Torrado, Philip Bontrager, Ahmed Khalifa, Julian Togelius, and Sebastian Risi. 2018. Illuminating Generalization in Deep Reinforcement Learning through Procedural Level Generation. *NeurIPS Workshop on Deep Reinforcement Learning* (2018).
- [17] Yoon Kim, Yacine Jernite, David Sontag, and Alexander M. Rush. 2016. Character-aware neural language models. In *AAAI'16 Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*. 2741–2749.
- [18] Jialin Liu, Sam Snodgrass, Ahmed Khalifa, Sebastian Risi, Georgios N Yannakakis, and Julian Togelius. 2020. Deep learning for procedural content generation. *Neural Computing and Applications* (2020), 1–19.
- [19] Tomas Mikolov, Kai Chen, Greg S. Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *ICLR (Workshop Poster)*.
- [20] Adam Roberts, Jesse Engel, Colin Raffel, Curtis Hawthorne, and Douglas Eck. 2018. A hierarchical latent vector model for learning long-term structure in music. *arXiv preprint arXiv:1803.05428* (2018).
- [21] Matthew Stephenson and Jochen Renz. 2017. Generating varied, stable and solvable levels for angry birds style physics games. In *2017 IEEE Conference on Computational Intelligence and Games (CIG)*. 288–295.
- [22] Matthew Stephenson and Jochen Renz. 2019. Agent-Based Adaptive Level Generation for Dynamic Difficulty Adjustment in Angry Birds. *AAAI-2019 Workshop on Games and Simulations for Artificial Intelligence* (2019). arXiv:arXiv preprint arXiv:1902.02518
- [23] Matthew Stephenson, Jochen Renz, Xiaoyu Ge, Lucas Ferreira, Julian Togelius, and Peng Zhang. 2019. The 2017 AIBIRDS Level Generation Competition. *IEEE Transactions on Games* 11, 3 (2019), 275–284.
- [24] Adam Summerville, Sam Snodgrass, Matthew Guzdial, Christoffer Holmgard, Amy K. Hoover, Aaron Isaksen, Andy Nealen, and Julian Togelius. 2018. Procedural Content Generation via Machine Learning (PCGML). *IEEE Transactions on Games* 10, 3 (2018), 257–270.
- [25] Sarjak Thakkar, Changxing Cao, Lifan Wang, Tae Jong Choi, and Julian Togelius. 2019. Autoencoder and Evolutionary Algorithm for Level Generation in Lode Runner. In *2019 IEEE Conference on Games (CoG)*. 1–4.
- [26] Ruben Rodriguez Torrado, Ahmed Khalifa, Michael Cerny Green, Niels Justesen, Sebastian Risi, and Julian Togelius. 2019. Bootstrapping Conditional GANs for Video Game Level Generation. *arXiv preprint arXiv:1910.01603* (2019).
- [27] Vanessa Volz, Jacob Schrum, Jialin Liu, Simon M. Lucas, Adam Smith, and Sebastian Risi. 2018. Evolving mario levels in the latent space of a deep convolutional generative adversarial network. In *Proceedings of the Genetic and Evolutionary Computation Conference on*. 221–228.
- [28] Adeel Zafar, Shahbaz Hassan, et al. 2019. Corpus for Angry Birds Level Generation. In *2019 2nd International Conference on Computing, Mathematics and Engineering Technologies (iCoMET)*. IEEE, 1–4.