

# Moonshine: Distilling Game Content Generators into Steerable Generative Models

**Yuhe Nie<sup>\*1</sup>, Michael Middleton<sup>\*1</sup>, Tim Merino<sup>1</sup> Nidhushan Kanagaraja<sup>1</sup> Ashutosh Kumar<sup>1</sup>  
Zhan Zhuang<sup>†2,3</sup> Julian Togelius<sup>†1</sup>**

<sup>1</sup>New York University Game Innovation Lab, Brooklyn, New York, USA

<sup>2</sup>Southern University of Science and Technology, Shenzhen, China

<sup>3</sup>City University of Hong Kong, Kowloon, Hong Kong SAR China

yn2273@nyu.edu, Michael.Middleton@nyu.edu, tm3477@nyu.edu, nk3755@nyu.edu, ak10514@nyu.edu,  
12250063@mail.sustech.edu.cn, Julian@Togelius.com

## Abstract

Procedural Content Generation via Machine Learning (PCGML) has enhanced game content creation, yet challenges in controllability and limited training data persist. This study addresses these issues by distilling a constructive PCG algorithm into a controllable PCGML model. We first generate a large amount of content with a constructive algorithm and label it using a Large Language Model (LLM). We use these synthetic labels to condition two PCGML models for content-specific generation, the Five-Dollar Model and the Discrete Diffusion Model. This neural network distillation process ensures that the generation aligns with the original algorithm while introducing controllability through plain text. We define this text-conditioned PCGML as a Text-to-game-Map (T2M) task, offering an alternative to prevalent text-to-image multi-modal tasks. We compare our distilled models with the baseline constructive algorithm. Our analysis of the variety, accuracy, and quality of our generation demonstrates the efficacy of distilling constructive methods into controllable text-conditioned PCGML models.

**Keywords:** Procedural Content Generation, Synthesis Generation, Multimodal Machine Learning, Text-to-game-Map Generation

## Datasets —

[huggingface.co/datasets/DolphinNie/dungeon-dataset](https://huggingface.co/datasets/DolphinNie/dungeon-dataset)

## Introduction

Procedural Content Generation (PCG) is crucial to many types of video games, including roguelikes and other games that are designed around infinite content variety. PCG is also important for empowering users and designers to create various types of game content. Many games feature generators for content such as quests, maps, levels, and items. These generators are built on different algorithms and underlying principles, including search, constraint satisfaction, machine learning, and hand-crafted rules.

An important concern in content generation is controllability. Many content generators feature this to some degree;

in its simplest form, controllability can be selecting a map size or number of enemies. While this basic level of control may be suitable for some applications, it can be limiting in others. On the opposite end of the spectrum is text-conditioned content generation. In this form, any idea that can be expressed through natural language can serve as guidance for a generator. The rise of Text-to-Image (T2I) models has demonstrated the value of such flexible means of control over content generation. While Text-to-Image capabilities have improved massively in recent years, these models are not well suited for many game-content generation tasks, where each game has a unique set of constraints.

Procedural Content Generation via Machine Learning (PCGML) (Summerville et al. 2018) uses machine learning models to *learn* from existing game content and generate diverse artifacts. However, a common challenge for PCGML methods is the lack of quality data for training. This challenge is exacerbated when training controllable generators, which require diverse and descriptive labels for the data. To enable text-conditioned PCGML, the demand for large, richly labeled datasets becomes even greater. Data labeling, requiring significant human labor, is often infeasible for most game developers.

To address this hurdle, we propose a novel approach that integrates traditional PCG algorithms with the capabilities of Large Language Models (LLMs) to generate large synthetic datasets for text-conditioned PCGML. Specifically, we leverage the PCG algorithm from the game Brogue (Walker 2019) to create a diverse collection of game maps and automatically generate descriptive labels for each, referenced from a small set of human examples. This synthetic dataset is then used to train two Text-to-game-Map (T2M) generative models: the Five-Dollar Model (Merino et al. 2023), which is a feed-forward model, and a Discrete Diffusion Model, which will be explained in detail in this paper. The outputs of these models are evaluated through both quantitative and qualitative metrics. Ultimately, our approach enables generative models to replicate the behavior of a black-box PCG algorithm using only its generated artifacts, while incorporating natural language text conditioning.

This approach can be viewed as a form of knowledge distillation — traditionally used to transfer the capabilities of

<sup>\*</sup>These authors contributed equally.

<sup>†</sup>Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

one neural network to another. However, instead of distilling between neural networks, we are transferring the behavior of an arbitrary algorithmic process into a neural network. While this adaptation of distillation lacks a formal theoretical foundation, it demonstrates promising results in practice. Inspired by its unconventional and experimental nature, we refer to the overall approach as *Moonshine*.

In summary, our contributions are:

1. We devise a strategy for creating extensive synthetic datasets for text-conditioned PCGML using traditional PCG methods combined with LLMs and few-shot human labels.
2. We frame text-conditioned PCGML as a Text-to-game-Map (T2M) task, analogous to the Text-to-Image (T2I) task, but operating in a discrete domain, and propose two training methodologies to address its unique challenges.
3. We provide a distillation process to convert any traditional PCG algorithms into controllable text-conditioned PCGML models effectively.

## Related Work

### Procedural Dungeon Generation

Text-to-game-Map (T2M) generation is a form of Procedural Dungeon Generation, where PCG algorithms are used to generate maps for e.g. dungeon crawler-type video game (Viana and Dos Santos 2019). Togelius et al. (2011) identified two primary approaches in the current landscape of dungeon generation: constructive algorithms and search-based algorithms (Togelius et al. 2011). Constructive algorithms directly generate dungeon levels through a variety of methods, such as Cellular Automata (Wolfram 1983) and Generative Grammars (Horrocks 2014). Search-based algorithms utilize metaheuristics, such as evolutionary algorithms (Raffe et al. 2014) and Wave Function Collapse (Nie et al. 2024) to optimize dungeon content against a set of criteria, using cycles of generation, evaluation, and selection to iteratively refine solutions. Human controllable content generation is important for designers to input their own preferences into generation. However, Viana and Dos Santos (2019) highlighted the scarce reliance on mixed-initiative approaches within dungeon generation, where human design interacts with computer-generated content.

In the context of PCG, van der Linden, Lopes, and Bidarra (2014) proposed a methodology for generating dungeons in role-playing games using a combination of graph theory and geometric algorithms. Their approach focuses on creating diverse and engaging gameplay experiences through the procedural generation of dungeon layouts, room connections, and game objects. The authors demonstrated the effectiveness of their approach through an evaluation of the generated dungeons, showcasing the potential of procedural content generation in enhancing gameplay variety and player engagement. They identified controllability as the key challenge—ensuring that procedurally generated dungeons adhere to the desired criteria for gameplay progression, difficulty curves, and thematic consistency.

### Synthetic Data Labeling

Most synthetic data labeling approaches come from the domain of text-to-image (T2I) generative models and captioning. He et al. (2024) introduced PRISM, an algorithm that automatically identifies human-interpretable and transferable prompts that can effectively generate desired concepts. This is particularly noteworthy as prompt engineering has been recognized as an effective method for controlling the output of T2I generative models, but it is also laborious due to the need for manually crafted prompts.

Expanding on the theme of automation in machine learning, Yang et al. (2023) introduced Adaptive Language-Image Pre-training (ALIP), a bi-path model that incorporates supervision from both raw text and synthetic captions. This approach addresses the challenges posed by intrinsic noise and mismatched image-text pairs in web data, which can negatively impact representation learning. The effectiveness of ALIP was demonstrated through experiments across various model scales and pre-training datasets, achieving state-of-the-art performance on multiple downstream tasks.

### Text-to-Image Generation

In the field of computer vision, Generative Adversarial Networks (GANs) (Reed et al. 2016) have been widely used for text-to-image generation, leveraging adversarial training to produce high-quality outputs. However, GANs often struggle with mode collapse, and are not always responsive to conditional inputs. Diffusion Models (Saharia et al. 2022) address these issues by iteratively denoising random noise to generate images, achieving superior fidelity and diversity. Key components in Diffusion Models include Variational Autoencoders (VAEs) (Kingma, Welling et al. 2019), which provide structured latent spaces for controlled generation, and the Contrastive Language-Image Pre-training (CLIP) (Radford et al. 2021), which maps images and text into a shared multimodal embedding space, consists of an image encoder for visual inputs and a text encoder for textual descriptions. These encoders are trained jointly using a contrastive loss, enabling CLIP to learn a wide range of visual concepts directly from natural language descriptions. This allows it to perform various vision tasks in a zero-shot manner by comparing the embeddings of text prompts with those of images to identify the best match.

### Synthetic Dataset Generation

Our pipeline for generating synthetic  $\{map, description\}$  dataset follows three main steps:

1. Extract a map data point from a game using a traditional PCG algorithm.
2. Perform metadata and heuristic analysis on the map.
3. Input metadata and heuristics into an LLM prompt to generate a descriptive map label.

### Map Extraction

We extract the map data from the open-source rogue-like dungeon game Brogue (Walker 2019). This game uses a traditional constructive PCG algorithm to generate varied

and complex dungeon maps. Unlike many constructive algorithms such as Binary Space Partitioning or Cellular Automata (Shaker, Togelius, and Nelson 2016), Brogue’s algorithm ensures *playability*, considering the connectivity, terrain layout, and ecosystem when generating various difficulties for players.

The original Brogue game contains diverse tilesets with terrain, objects, monsters, items, and more. This research focuses on the terrain tiles to simplify the original game content. We scale the maps to a consistent size  $32 \times 32$  pixels with a tileset of 14 terrain tiles, as detailed in Table 1.

Tile	Desc.	Tile	Desc.	Tile	Desc.
	Ashes		Bog		Bridge
	Crystal		Fire		Fungus
	Grass		Ground		Ice
	Lava		None		Sand
	Stone		Water		

Table 1: The details of dungeon tileset with descriptions.

We create a dataset of maps from the game, split into 49,000 training points, 14,000 test points, and 7,000 validation points. We open-sourced this dataset and uploaded it to Hugging Face’s dataset repository.

## Map Metadata Analysis

We perform multiple heuristic calculations to extract metadata for each map. The metadata will then be provided in the LLM’s prompt to generate a descriptive label for the map. A visual example of metadata is shown in the Fig. 1.

1. We place a binary mask over each independent room and connecting path.
2. We divide the map into a grid of cardinal and inter-cardinal directions (N, S, E, W, NE, NW, SE, SW), and assign each room a direction using its midpoint.
3. We extract the tile counts for each room, order them by quantity, and add them to the room’s description.
4. We record the connected room pairs and the path to connect the pair.

## Synthetic Description Generation

**Pre-generation Prompt** We structure our pre-generation prompt into four key sections: The *Setting* introduces the task and clarifies the LLM’s goal. The *Response* format specifies the exact output structure required. The *Examples* provides a few-shot demonstration of human-authored map descriptions to guide the LLM. Finally, the *Rules* outlines specific constraints that the LLM must follow during generation.

**Generation Round** In each following round of conversation, we provide the LLM with an integer grid representation of the map, a dictionary of integers map to the tile names, and the meta-data information. We task the LLM to generate 10 text descriptions for each map. This process is repeated for each unique map.

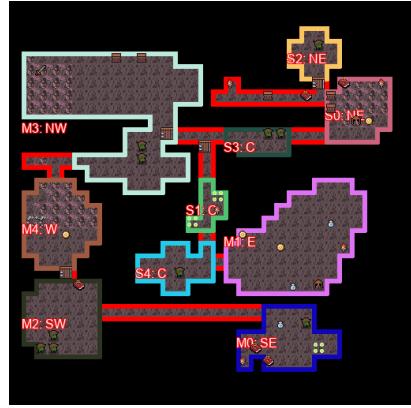


Figure 1: A visualized example of map metadata information. Independent rooms and connecting paths are colored. Room labels and directions are overlayed as text.

## Text-to-game-Map Generation

### Task Definition

Our approach to Text-to-game-Map (T2M) generation is similar to Text-to-Image generation but differs due to the discrete grid nature of the map data. Rather than generating continuous pixel values, our model generates a grid of probability vectors and then selects the most probable tile at each cell.

**Map Representation** Each map  $m$  in our dataset is represented as a three-dimensional matrix of dimensions  $H \times W \times C$  where  $H$  and  $W$  denote the map’s height and width, respectively, and  $C$  represents the tileset size. Each cell  $m_{i,j}$  within the matrix is a vector of length  $C$  and follows a classification distribution, ensuring that the sum of probabilities across all tiles in a cell equals one.

In the training dataset, each map is represented in a one-hot encoded format, where each cell is set to indicate a specific tile choice  $m_{i,j,k^*} = 1$ , while all other tiles are set to zero,  $m_{i,j,l} = 0$  for  $l \neq k^*$ . Mathematically, this is expressed as:

$$\forall i \in [0, H]; j \in [0, W]; \sum_{k=1}^C m_{i,j,k} = 1.$$

**Text Embedding Model** To extract text embedding vectors, we employ the pre-trained model *gte-large-en-v1.5* (Li et al. 2023), which is one of the best-performing small models on the massive text embedding benchmark. This model supports a maximum input size of 8192 tokens and generates 1024-dimensional embedding vectors.

**Generative Model** We define a comprehensive generative model  $\mathcal{G}$ , which produces maps  $M$  conditioned on text embedding  $\mathbf{t}$ . This model, parameterized by weights  $\theta$ , may optionally incorporate input  $\mathbf{z}$  from a noise distribution to enhance generation diversity. By integrating textual context into the map generation process, the model produces content that is both contextually relevant and varied, based on the input text.

We evaluate two multi-modal strategies to investigate the effects of distilling the constructive PCG algorithm into controllable T2M models.

### Five-Dollar Model (FDM)

The Five-Dollar Model (FDM), proposed by Merino et al. (2023), is a streamlined model for Text-to-game-Map generation. It is a feed-forward neural network, which maps a text embedding vector directly to a map output. The transformation from text to map can be represented as:

$$M' = \mathcal{G}_\theta(\mathbf{T}).$$

The model is only designed to minimize the reconstruction error via self-supervised learning as:

$$L = \|M - M'\|_2^2.$$

In this strategy, training can be conceptualized as a classification task. The text embedding vector serves as a discriminative feature that the model uses to "classify" into the correct map representation. This perspective aligns the generative task with classification paradigms, focusing on feature utilization and output accuracy.

Fig. 2 outlines the network architecture. Basic feedforward layers are employed to align the text embedding vector to the map. The text embedding vector is concatenated with a small noise vector, passed through a dense layer, and reshaped into a grid. This is upsampling through three residual blocks, and finally reshaped to a  $H \times W \times C$  map output via a convolutional layer.

### Discrete Diffusion Model (DDM)

To tackle the diversity challenge, we provide one specific method for training a Discrete Diffusion Model (DDM). It starts with a  $H \times W \times C$  array of random normal noise map  $m_t$  and gradually denoises it across multiple timesteps, each time moving closer to the map distribution. This iterative refinement uses a model that learns and predicts the noise to be removed at each step (Rombach et al. 2022). Conditioning on the text embedding vector ensures that the denoising trajectory is aligned with the semantics of the input text, allowing the generation of maps that possess visual appeal and are contextually relevant to the input text.

As illustrated in Fig. 3, our diffusion model  $\mathcal{G}_\theta$  is based on a conditional UNet comprising Resnet and cross-attention blocks that effectively learn the constrained content passed from the text embeddings. At each time step  $t$ , the model takes a noisy map  $m_t$  and the text embedding vector  $\mathbf{t}$  as input, and outputs the predicted noise  $\hat{\epsilon}$  of the map. We adopt the same setting as the Denoising Diffusion Probabilistic Model (DDPM) (Ho, Jain, and Abbeel 2020) to add the Gaussian noise to the map at each timestep. The denoising process can be mathematically formulated as follows:

$$m_t = \text{DDPM}(\epsilon, m_0, t) = \sqrt{\bar{\alpha}_t} m_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \\ \text{where } \epsilon \sim \mathcal{N}(0, 1) \text{ and } \hat{\epsilon} = \mathcal{G}_\theta(m_t, \mathbf{t}, t).$$

In the DDPM,  $\bar{\alpha}_t$  denotes a decreasing scheduler. The training objective involves minimizing a loss function which

focuses on the accurate prediction of noise,  $\epsilon$ . The loss function,  $L$ , is defined as follows:

$$L = \mathbb{E}_{\mathbf{m}_0, \mathbf{t}, \epsilon \sim \mathcal{N}(0, 1), t \sim \mathcal{U}(1, T)} (\|\epsilon - \mathcal{G}_\theta(\mathbf{m}_t, \mathbf{t}, t)\|_2^2).$$

This iterative denoising process continues until  $t = 0$ , resulting in a map  $m_0$  that resembles the original training data while incorporating the semantics of input text  $\mathbf{t}$ .

## Experiment and Result

### Evaluation Methods

We use OpenAI's GPT-4 Turbo (*gpt-4-turbo-2024-04-09*) to generate map descriptions. Due to financial constraints, we only generate descriptions for 3,000 maps (in principle, an arbitrary number of descriptions could be generated). Each map is given 10 LLM-generated descriptions, divided into 5 long and 5 short descriptions. Long and short descriptions vary in detail and breadth. For example, the long and short descriptions for one map are:

- **Long Description:** *A diverse terrain with four main areas, each featuring a combination of fungus and ground. The northwest region is dotted with stone and ashes amidst more ground and fungus.*
- **Short Description:** *Four area division: ground, fungus, scarce stones, and ash fragments.*

We evaluate the descriptions using the following metrics:

- **BLEU** (Papineni et al. 2002) and **ROUGE-L** (Lin 2004) scores assess the overlap between generated descriptions and human written references, indicating how closely the generated content resembles human language in terms of lexical choice and phrasing. Higher scores suggest that the model produces text that is more similar to human descriptions, reflecting better human-like quality.
- **METEOR** (Lavie and Agarwal 2007) considers synonymy and stem matching, providing insight into the semantic adequacy of the descriptions. Higher scores suggest that the content conveys the same meaning as human descriptions.
- **SPICE** (Anderson et al. 2016) focuses on the semantic propositional content, evaluating how well the generated descriptions capture the underlying meaning and relationships present in human descriptions. This metric is particularly relevant for assessing the richness of semantic content and the diversity of generated outputs.
- **CLIP Score** (Hessel et al. 2022): Measures *cross-modal alignment* between text and images using a pretrained vision-language model.

These metrics collectively assess the descriptions for lexical overlap, semantic content, and alignment with human-like language, ensuring diversity and contextual relevance.

### Evaluation of Generated Descriptions

We separately evaluate the effects of the generated long and short descriptions, including comparisons within each type and against human references.<sup>1</sup>

<sup>1</sup> Additional comparisons can be found in the appendix.

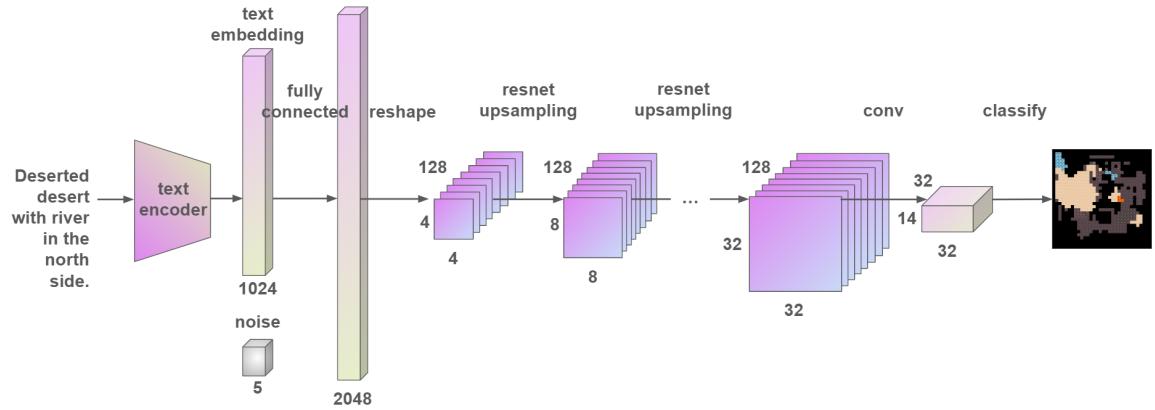


Figure 2: Five-Dollar Model architecture.

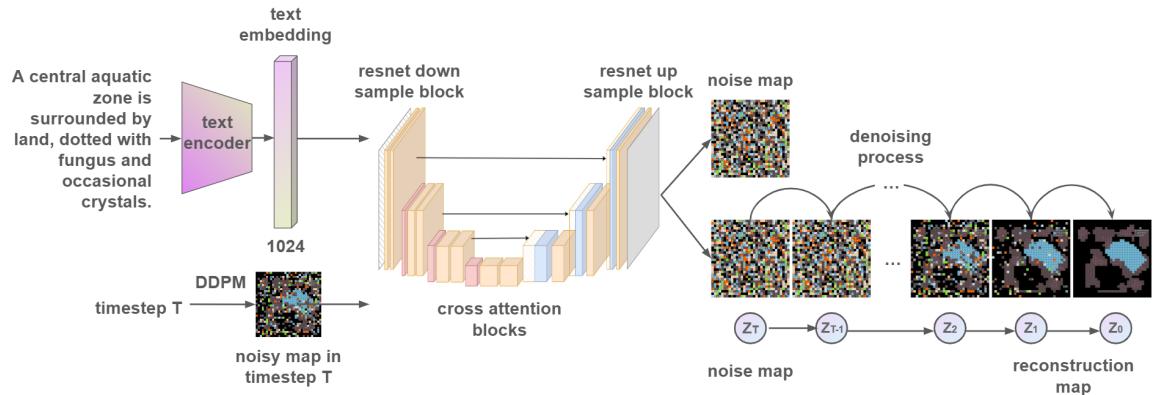


Figure 3: Discrete Diffusion Model architecture.

S.No	Metric	Des. Long	Des. Short
<b>0</b>	Bleu_1	54.71	<b>58.37</b>
<b>1</b>	Bleu_2	<b>26.57</b>	26.56
<b>2</b>	Bleu_3	<b>12.18</b>	11.68
<b>3</b>	Bleu_4	<b>05.62</b>	04.86
<b>4</b>	METEOR	<b>19.50</b>	12.64
<b>5</b>	ROUGE_L	<b>33.27</b>	26.16
<b>6</b>	SPICE	<b>11.31</b>	03.89

Table 2: Comparison within generated descriptions.

**Comparison within Generated Descriptions** For each type of description, we selected the first description as a reference, calculated the scores against the other four descriptions, and averaged the results across all 3,000 maps (Table 2). Long descriptions outperformed short descriptions in almost all metrics, suggesting that long descriptions better align with human-like quality and capture more detailed and relevant information.

**Generated Descriptions versus Human References** We compared synthetic captions with 30 randomly selected human references to evaluate how closely synthetic captions

S.No	Metric	Des. Long	Des. Short
<b>0</b>	Bleu_1	<b>70.84</b>	50.58
<b>1</b>	Bleu_2	<b>43.47</b>	21.13
<b>2</b>	Bleu_3	<b>23.56</b>	07.88
<b>3</b>	Bleu_4	<b>11.49</b>	02.38
<b>4</b>	METEOR	<b>16.56</b>	09.38
<b>5</b>	ROUGE_L	<b>42.92</b>	23.25
<b>6</b>	SPICE	<b>03.43</b>	00.77

Table 3: Comparison between generated descriptions and human references.

resemble human descriptions (Table 3). These results are similar to the previous comparison. Long descriptions more effectively mimic the human descriptive style than short descriptions, likely because they incorporate more detail, making them appear more natural and comprehensive.

The use of random human references as a benchmark is insightful, revealing that long synthetic descriptions maintain strong linguistic and semantic quality, even when compared against diverse human-written descriptions.

Brogue	FDM	DDM

Table 4: Generated maps based on generated descriptions: (1) A **central area dominated by a large expanse of lava** surrounded by **solid ground**, with **sparsely growing fungi** and **a few patches of grass**; (2) A **vast desert** landscape merges with **clear blue waters**, spotted with **fungus** growths and **small patches of greenery** amidst the **sandy terrain**; (3) A **central aquatic zone** is surrounded by **land**, **dotted with fungus** and **occasional crystals**. Two smaller regions lie on the **eastern edge**, connected by pathways to the main area; (4) A **charred landscape** with **scorched earth** and **rampant wildfires** spread across. **Patches of fungi** cling to life amidst the destruction, with rare **mineral-rich stones** scattered **sparsely**.

## Evaluation of Generated Maps

Tables 4 and 5 show generated maps using the Moonshine distillation process. Table 4 compares source maps from Brogue with generated maps, using LLM-generated descriptions. Table 5 compares source and generated maps using human-authored input. Specific tile types, locations, and ranges in prompts are highlighted in **red**, **blue**, and **orange**, respectively. We see that DDM can produce a diverse set of maps compared to FDM. However, under a close evaluation FDM sometimes captures nuances of the text slightly better. The FDM sand tiles cover nearly the entire map in "A vast sandy area", whereas the sand is a smaller portion in the DDM generations. When asked to create "some lakes", FDM produces multiple lakes whereas DDM produces only one lake. As the diffusion method relies on data, it is possible that with more examples given, the model may start to understand these smaller nuances in the prompts.

FDM	DDM-1	DDM-2

Table 5: Generated maps based on human prompts: (1) A **vast sandy** area; (2) Some **lakes** in the **north**, and a **lot of magma** and **lava**; (3) **Burning out fire** in the **center**; (4) The **lake** to the **left**. The **desert** to the **right**. Connected by **bridges**.

**Quantitative Evaluation** The models demonstrate reasonable accuracy in reflecting specific types of tiles mentioned in the descriptions, especially for the main area. Both models excel at capturing descriptions like "*central area*" or "*a vast area of...*" by generating wide regions that include the specified tiles. However, they sometimes struggle with more detailed elements, such as "*a dot of...*", "*scattered around*," or "*few...*" This limitation may stem from synthetic descriptions that primarily focus on the main area while neglecting finer details.

While not all generated regions are connected, the models exhibit an understanding of hidden *connectivity* and *ecology* features beyond explicit descriptions. FDM struggles with diversity, producing fixed outputs based on a given prompt due to its focus on learning mapping relationships. In contrast, DDM outperforms it by generating *varied* results that align with the text and include additional details.

**Similarity between Generated Map and Ground Truth** We calculate the CLIP Score between the ground truth maps Brogue generated and the maps generated by the FDM or DDM, respectively. The box plots are shown in Fig. 4 which exhibits a concentrated distribution of data points, showing a strong correlation of the similarity between the Brogue ground truth and the predictions of the two generative models. DDM demonstrates a high degree of clustering, indicating good predictive accuracy.

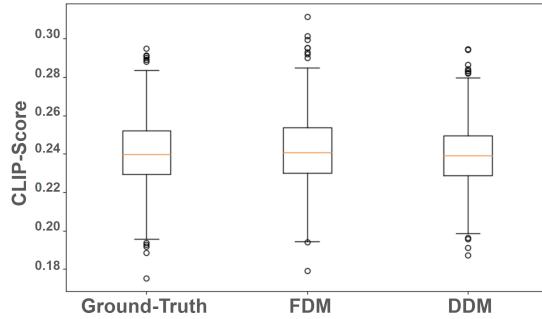


Figure 4: CLIP score of FDM (left) and DDM (right) v.s. the Brogue ground truth (X-axis).

To compare the diversity of DDM and the FDM, we perform connected component analyses to assess usable structural diversity. The results in Fig. 5 demonstrate that the DDM generates maps with a wider range of usable patterns.

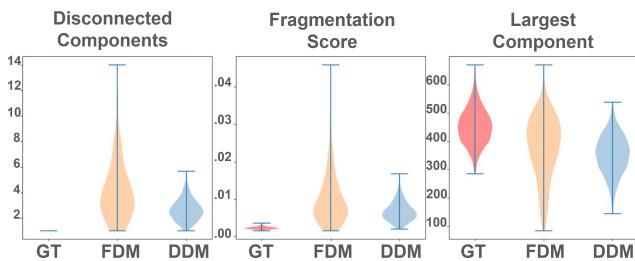


Figure 5: Connectivity analysis of models based on three metrics: the number of disconnected components, fragmentation score, and largest component size. DDM has fewer disconnected components, lower fragmentation, and more stable component sizes compared to FDM.

## Conclusion

We introduce Moonshine, a method for distilling traditional PCGML algorithms into text-conditioned PGCML models using synthetic data generation. Based on an automated map-labeling approach via Large Language Models (LLM), we minimize the labor for human annotation — a major obstacle for Text-to-game-Map (T2M) models. Additionally, we provide our open-source dataset of dungeon maps with associated descriptions for use in further T2M research.

By combining traditional constructive PCG algorithms with the semantic understanding of LLMs, we enable the creation of diverse, human-like synthetic descriptions for existing PCG map data. We found that longer descriptions generated by LLMs capture more semantic information and align more closely with human references, whereas shorter descriptions correspond more with map image information in pretrained CLIP models. However, our fine-tuned CLIP model shows that longer descriptions integrate better with maps, suggesting that while a frozen model may perform better with short descriptions, longer descriptions are preferable for fine-tuning multimodal-aligned models.

We explored two T2M models —the Five-Dollar Model (FDM) and the Discrete Diffusion Model (DDM) — and found that while both generate maps correlated with their respective text descriptions, longer descriptions offer more accurate semantic reflections. FDM utilizes a feed-forward network to directly map text embeddings onto a map, achieving slightly higher overall scores compared to the DDM. However, it lacked diversity, producing relatively uniform results. In contrast, the DDM, which incrementally restores a noise map to a reconstructed map, shows greater potential for diversity.

However, the limitations in the dataset size for training the DDM suggest that extending training over more epochs with a larger synthetic dataset could improve its performance. We acknowledge that a loss of fine-grained control is a limitation of our current model, partly due to constraints in LLM annotations and the available training data. We anticipate that expanding the dataset to include additional tile types could mitigate this issue and enhance the model’s level of control. Additionally, A larger dataset of human-written descriptions would provide valuable information regarding the human likeliness of the generated descriptions. A study should be conducted to assess the use of the distilled Moonshine models with a designer trying to construct specific map outputs. We look only at a subset of procedural dungeon generation in terrain: items, enemies, and even stories all contribute to better model distillation.

In summary, Moonshine’s automated distillation generation process can be an effective method for creating controllable content generators. For optimal results in description generation and model selection, we recommend prioritizing longer descriptions and considering the Discrete Diffusion Model for tasks requiring diverse outputs. Future research can explore ways to encode additional metadata into the map descriptions, such as items, enemies, and stories, or incorporate reinforcement learning (RL) mechanisms to further refine control over generated maps, allowing dynamic user interaction during generation.

## Appendix

All supplemental materials are available at <https://arxiv.org/abs/2408.09594> in section titled "Appendix", released under a CC BY 4.0 license. They include (1) Pre-trained v.s. Fine-tuned CLIP Model for Synthetic Descriptions, (2) Diversity of the Synthetic Descriptions, (3) Visualization of DDM, (4) Reconstruction Loss (5) Prompt Formation. The example prompts and the full dataset can be found at [huggingface.co/datasets/DolphinNie/dungeon-dataset](https://huggingface.co/datasets/DolphinNie/dungeon-dataset)

## References

Anderson, P.; Fernando, B.; Johnson, M.; and Gould, S. 2016. SPICE: Semantic propositional image caption evaluation. In *European Conference on Computer Vision (ECCV)*, 382–398.

He, Y.; Robey, A.; Murata, N.; Jiang, Y.; Williams, J.; Pappas, G. J.; Hassani, H.; Mitsufuji, Y.; Salakhutdinov, R.; and Kolter, J. Z. 2024. Automated Black-box Prompt

Engineering for Personalized Text-to-Image Generation. *arXiv*:2403.19103.

Hessel, J.; Holtzman, A.; Forbes, M.; Bras, R. L.; and Choi, Y. 2022. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. *arXiv*:2104.08718.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.

Horrocks, G. 2014. *Generative grammar*. Routledge.

Kingma, D. P.; Welling, M.; et al. 2019. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, 12(4): 307–392.

Lavie, A.; and Agarwal, A. 2007. METEOR: An automatic metric for MT evaluation with high levels of correlation with human judgments. In *Proceedings of the ACL Workshop on Statistical Machine Translation*, 228–231.

Li, Z.; Zhang, X.; Zhang, Y.; Long, D.; Xie, P.; and Zhang, M. 2023. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv*:2308.03281.

Lin, C.-Y. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out, Post-Conference Workshop of ACL*, 74–81.

Merino, T.; Negri, R.; Rajesh, D.; Charity, M.; and Togelius, J. 2023. The five-dollar model: generating game maps and sprites from sentence embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 19, 107–115.

Nie, Y.; Zheng, S.; Zhuang, Z.; and Togelius, J. 2024. Nested Wave Function Collapse Enables Large-Scale Content Generation. *IEEE Transactions on Games*.

Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics (ACL)*, 311–318.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.

Raffe, W. L.; Zambetta, F.; Li, X.; and Stanley, K. O. 2014. Integrated approach to personalized procedural map generation using evolutionary algorithms. *IEEE Transactions on Computational Intelligence and AI in Games*, 7(2): 139–155.

Reed, S.; Akata, Z.; Yan, X.; Logeswaran, L.; Schiele, B.; and Lee, H. 2016. Generative adversarial text to image synthesis. In *International conference on machine learning*, 1060–1069. PMLR.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.

Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35: 36479–36494.

Shaker, N.; Togelius, J.; and Nelson, M. J. 2016. Procedural content generation in games. *Procedural content generation in games*.

Summerville, A.; Snodgrass, S.; Guzdial, M.; Holmgård, C.; Hoover, A. K.; Isaksen, A.; Nealen, A.; and Togelius, J. 2018. Procedural content generation via machine learning (PCGML). *IEEE Transactions on Games*, 10(3): 257–270.

Togelius, J.; Yannakakis, G. N.; Stanley, K. O.; and Browne, C. 2011. Search-based procedural content generation: A taxonomy and survey. *IEEE Transactions on Computational Intelligence and AI in Games*, 3(3): 172–186.

van der Linden, R.; Lopes, R.; and Bidarra, R. 2014. Procedural Generation of Dungeons. *IEEE Transactions on Computational Intelligence and AI in Games*, 6(1): 78–89.

Viana, B. M. F.; and Dos Santos, S. R. 2019. A Survey of Procedural Dungeon Generation. In *2019 18th Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*, 29–38. Rio de Janeiro, Brazil: IEEE. ISBN 978-1-72814-637-9.

Walker, B. 2019. Brogue: Community edition - A community-lead fork of the much-loved minimalist rogue-like game.

Wolfram, S. 1983. Statistical mechanics of cellular automata. *Reviews of modern physics*, 55(3): 601.

Yang, K.; Deng, J.; An, X.; Li, J.; Feng, Z.; Guo, J.; Yang, J.; and Liu, T. 2023. ALIP: Adaptive Language-Image Pre-Training with Synthetic Caption. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2922–2931.