Proceedings of CONF-SEML 2026 Symposium: Computational Analysis and Modeling in Complex Intelligent Systems
DOI: 10.54254/2755-2721/2026.GU34636

Evolution of AI-Driven Game Space Generation: Deep Learning, Reinforcement Learning, and Neuro-Symbolic Architectures

Yushan Chen

Department of Aeronautical and Aviation Engineering, The Hong Kong Polytechnic University, Hong Kong, China
yu-shan.chen@connect.polyu.hk

Abstract. In non-combative games (such as meat pigeon, life simulation and box puzzle games), the quality of architectural space directly affects the player's exploration process of the scene and the effect of environmental narrative. Traditional procedural content generation (PCG) relies heavily on heuristic search and rigid rules, and it is difficult to balance the rationality of "Space Syntax" and architectural aesthetics in complex architectural topology. This article systematically summarises the evolution of game architecture PCG technology based on deep learning and reinforcement learning in the past five years. Firstly, it sorts out the history of technological development, and secondly, it evaluates different technical schools from three dimensions: high-dimensional non-convex feature space optimisation, semantic rule quantification via reward shaping mechanisms, and closed-loop topological quality verification based on Space Syntax. Through horizontal comparison, this article points out that single generation algorithms often have inherent limitations and compromises, such as the topological vulnerability of the generation model and the aesthetic rigidity of heuristic rules. Then it predicts the future development direction, emphasising the potential of neural symbol architecture and large-scale language models (LLMs) in realising human-like spatial narrative.

Keywords: Procedural Content Generation, Game Architecture, Deep Reinforcement Learning, Space Syntax

1. Introduction

With the rapid development of game industrialisation, the positioning of game scenes in the complex system of games is also constantly changing. The game scene is not only a background with artistic significance, but also a dynamic system that carries the core interaction mechanism, emotional narration and even passive environmental narrative. Especially in the game types that emphasise exploration and experience (such as Japanese Role-Playing Game (JRPG), life simulation, box puzzle solving, etc.), the game scene design itself is one of the core drivers of the gameplay.

For non-COMBAT games, the fun of players often comes from the desire to explore unknown spaces, curiosity about stories, the emotional value of interacting with Non-Player Characters

DOI: 10.54254/2755-2721/2026.GU34636

(NPCs) in specific scenes, or the sense of accomplishment of solving environmental puzzles. The above contents put forward great requirements for the generation of the environment and buildings. Game architecture should not only conform to the streamlined logic of real architecture, but also meet the rational needs of functional design in the game, and also need other emotional added values. An excellent non-combative game space should be able to perfectly balance "provide players with the freshness of multi-week play" and "in line with design logic". Procedural content generation (PCG) technology is the key technology that can reduce production time and cost while organically balancing the above needs.

In early game development, the automated generation of space mainly relies on some classic algorithms (such as A* pathfinding, cell automaton or binary space division (BSP) tree). For example, the BSP tree can divide the available space into smaller rectangular areas and then connect them, which can efficiently generate a basic interior layout of the dungeon that does not have complex design requirements. The biggest advantage of this kind of method, based on propositional logic and geometric heuristics, is the speed and certainty of the calculation, which can quickly generate a connected basic space.

However, with the improvement of players' aesthetic level and requirements for the game, the traditional method has gradually reached a bottleneck. First, the method of exhaustive rules is prone to face homogenisation in the face of three-dimensional architecture with high freedom and multiple constraints: three-dimensional design needs to face exponential growth, and it is difficult for the algorithm to exhaust all legal methods; multi-constraints may make the algorithm to face deadlock, and at this time the algorithm will choose to push it back mechanically instead of Improvise like human beings. Levels generated by pure algorithms often lack the depth of spatial configuration and are difficult to induce complex player behaviour, leading to strong spatial homogenisation and mechanical sense [1]. Secondly, it is difficult for traditional algorithms to deal with the highly complex characteristics of "narrative-space binding". As pointed out by the research of Lech et al. [2], when the game needs to generate coherent world background, character tasks and corresponding specific physical buildings at the same time, the traditional single generation assembly line often falls into a logical dead end due to too many constraints, resulting in a serious disconnect between gameplay and narrative.

In order to go further on the traditional geometric rules, in recent years, research has changed from traditional deterministic logic algorithms to artificial intelligence (AI) machine learning. As pointed out by Snodgrass et al. [3], the research focus of academia and industry has shifted on a large scale to machine learning-based program-based content generation (PCGML). This change has brought about two technological leaps. First, the generative AI represented by Convolutional Neural Networks (CNN), Generative Adversarial Networks (GAN) and diffusion models greatly improves the coherence and style diversity of game scene design by learning a large number of game scenes and level data designed by excellent human designers; second, deep learning has completely changed the underlying logic of content generation, and the static way-finding Algorithm and spatial planning turn into a dynamic trial and error interaction process of agent to pursue reward maximisation in the virtual environment. Against this background, this article aims to comprehensively summarise this technological progress, deeply analyse the processing of game architecture space topology by different machine learning algorithms, compare their mechanical advantages and limitations under multi-modal constraints, and further think about the possible development path of game architecture generation in the future.

DOI: 10.54254/2755-2721/2026.GU34636

2. Current research status and technical classification

According to several different aspects such as the driving mechanism at the bottom, state space characterisation (how AI views and understands the game environment), and how to connect the topological space, the PCG technology of game architecture space has continued to evolve in the past five years. This paper can divide the current mainstream technology into the following four core stages:

2.1. Classic heuristics and hard rules

Before the widespread popularisation of generative AI and deep learning, people tended to use traditional algorithms to generate game architecture space. Such algorithms tend to pursue absolute physical connections and geometric non-overlapping. Representative techniques at this stage include graph theory-based cutting (such as BSP tree), cell automatons, and wave function collapse (WFC) algorithms. Although this kind of algorithm has very little computing overhead and can ensure the basic geometric rationality of space (connected and not overlapping), it cannot meet the increasingly complex aesthetic needs of non-combat games and the in-depth needs of level mechanism design. At this stage, Space Syntax also began to be introduced for preliminary graphical evaluation of the generated spatial topology [1].

2.2. PCGML

In order to break through the generation limitations brought about by the traditional manual writing of hard rules, the academic community began to introduce PCGML based on machine learning. In the early days, PCGML adopted a large number of CNN and GAN, trying to use the unique local perception and parameter sharing mechanisms of these models to automatically learn and reproduce the structural topology and aesthetic laws of the game building space from a large amount of high-quality artificial design data.

In recent years, the Unconditional Diffusion model has been successfully applied to the generation of discrete levels [4]. In Addition, In Response To The Common Data Scarcity Problem In Game Development, Dai et al. [5] Proposed A Diffusion Model Framework Based On Single Example. By Introducing Intensive Sementic Labels And Restricted Sensing Fields, It Realises That More Can Be Generated With Only A Single Architectural Reference Map. The breakthrough in the layout of the variety.

2.3. PCGRL

Although the generative large model is realistic, there are still certain defects in logic. For this reason, Khalifa et al. [6] proposed the pioneering Procedural Content Generation via Reinforcement Learning (PCGRL) framework. They proposed that generation is no longer regarded as a one-time image output, but modelled as a Markov decision-making process (MDP). In the random initial state, the agent is designed to repeatedly perform spatial construction actions such as placing walls and doorways. On this basis, Kerker et al. [7] further introduced the Reward Shaping mechanism, compiling complex game-building rules into environmental rewards, and guiding the intelligent body to learn the correct spatial semantics.

DOI: 10.54254/2755-2721/2026.GU34636

2.4. Large language models and neural-symbolic architectures

With the outbreak of the Large Language Model (LLM), generation technology has ushered in a new era. LLMs have strong common-sense reasoning ability and can understand complex spatial languages. In order to overcome the problems of "hallucination" and state explosion in the calculation of pure large models in three-dimensional space, the latest research proposes a "neurological symbol architecture". For example, Xu and Verbrugge [8] designed a 3D generation framework based on LLMs and database-driven, using large models to generate high-level topological logic diagrams, and combining discrete 3D databases for underlying physical instantiation assembly.

3. Core methodology and optimization mechanism analysis

In modern research, the effectiveness of evaluating algorithms is no longer limited to operating efficiency, but needs to consider the geometric rationality, spatial semantics and other comprehensive capabilities of the space generated under multimodal constraints. In order to objectively compare the performance differences of the different underlying algorithms mentioned above, this section systematically sorts out the core methodology and evaluation benchmarks of game building space generation from the three dimensions of feature space optimisation, semantic rule quantification and topological quality verification.

3.1. The state explosion of spatial feature dimensions and non-convex optimization

The three-dimensional non-combative game building space contains a large number of heterogeneous physical features with great differences in dynamic range, such as discrete node connections, continuous wall thicknesses and the absolute area of the room. When these architectural attributes are converted into the input of the machine learning model, the dimension of the feature space grows exponentially.

From the perspective of numerical optimisation, this huge and uneven characteristic difference leads to a highly non-convex loss function space. In the process of gradient reverse propagation, the Condition Number of the Hessian matrix is very easy to increase, which is manifested as the violent oscillation of the optimisation path in the narrow loss manifold. Therefore, the modern generation model must build a sophisticated Feature Normalisation mechanism at the input layer to map high-dimensional heterogeneous properties into a smooth potential space. Whether the model has the geometric robustness to deal with high-dimensional state explosions and effectively converges to avoid physical collapses such as "piece rooms," is the primary criterion for evaluating its underlying computing power [9].

3.2. Semantic constraints and reward shaping mechanism in architecture

In life simulation and non-combat Role-Playing Game (RPG) games, the spatial layout not only needs to meet the requirements of physical anti-wearing, but also needs to meet the common sense expectations of human architecture (such as "the kitchen is adjacent to the dining room" or "the bathroom has a certain depth of privacy"). When dealing with such complex logic, the traditional reinforcement learning (RL) framework often falls into blind search because the reward signals are too sparse.

DOI: 10.54254/2755-2721/2026.GU34636

In order to realise the leap from "graphic connection" to "semantic correctness", Kerker et al. [7] proposed the use of a rigorous Reward Shaping mechanism. The mechanism transforms discrete architectural design rules into differentiable dense mathematical feedback. Specifically, the method uses heuristic signals to guide the intermediate action close to the semantics of the target (such as opening a specific functional partition). At the same time, it uses domain knowledge to build a potential-based punishment function to forcibly cut off the generated action that violates common sense. Whether it can cross the semantic gap through reward plasticity or similar mechanisms and ensure the rationality of the common sense of generating space is the core indicator to measure the depth of understanding of spatial semantics by generating algorithms [10].

3.3. Topological quality closed-loop verification based on spatial syntax

Of course, the most important verification is actually the actual play experience of the game and the rationality of the player's dynamic line. The early generative model mainly focused on the fitting of visual features and lacked a graphical understanding of the connection between macroscopic starting and stopping points. In order to get rid of subjective and expensive playtesting, the existing optimisation framework gradually incorporates the "Space Syntax" theory into the quantitative evaluation system.

By calculating indicators such as "Integration" and "Choice" of architectural topology, spatial syntax can objectively evaluate the convenience of a specific spatial node as a transportation hub and the probability of being a must. In the evolution of modern PCG engines, the thermal map evaluation mechanism of spatial syntax is directly embedded in the loss function or verification module, so that the model can identify and punish the generation network with a high loss rate or inefficient motion lines in real time without calling the physical engine. This closed-loop verification ability that integrates topology self-examination effectively tests the rationality of the player's expected dynamic line [11].

4. Horizontal comparison and mechanism game theory

In non-combative games, the generation of architectural space is a multi-modal constraint problem that needs to meet the physical geometry of the bottom, the middle, topological connection and the aesthetics of high-rise buildings at the same time. Looking back on the four major algorithm schools described in the second section, this paper can find that they are limited by the underlying driving mechanism and show completely different advantages and limitations when dealing with these conflicting constraints.

4.1. Heuristic algorithm

The mechanical advantage of hard rule algorithms represented by BSP and WFC is that they can provide absolute physical guarantees (high connection rate). Because the algorithm strictly follows the preset anti-overlap rules, it can effectively avoid the phenomenon of mould wear. However, the shortcomings of its mechanism are also extremely obvious: extreme lack of aesthetics and flexibility (very low aesthetic fit). In the face of three-dimensional architecture with a high degree of freedom, the exhaustion of rules can very easily cause conflicts between constraints, so that the algorithm can only generate the most conservative "square box" layout in order not to make mistakes. This mechanism of sacrificing design white space for the sake of connectivity is the fundamental reason

DOI: 10.54254/2755-2721/2026.GU34636

for the serious homogenisation of traditional generation space and the difficulty of carrying complex environmental narratives.

4.2. PCGML

PCGML, represented by CNN and GAN, has the opposite mechanism performance to heuristic algorithms. With the help of strong feature extraction ability, this kind of model can perfectly imitate the implicit aesthetic laws of human designers (the aesthetic fit is extremely high). However, due to the local sensing field characteristics of convolutional networks, they essentially lack an understanding of the path of global graph theory. This means that AI can only see whether the partial room splicing is natural, but cannot guarantee whether there is an accessible path from the starting point to the end point at the macroscopic level, which leads to an extremely low spatial connection rate.

4.3. PCGRL

The PCGRL framework regains control of the spatial topology through dynamic trial and error of state-action. As long as the spatial rules of path-finding accessibility and common sense are translated into semantic punishment and written to the reward function, the intelligent body can force these logics to be satisfied in the MDP loop. But the bottleneck of RL lies in its highly utilitarian optimisation nature. Without extremely complex and expensive reward function fine-tuning (Reward Engineering), the intelligent experience will generate a visually unnatural torturous maze in order to pursue high scores, so it is usually only at a medium level in the performance of architectural aesthetics.

4.4. Neural-symbolic hybrid architecture

In view of the limitations of a single mechanism, the "LLM + 3D database" architecture shows dynamic balance under multi-modal constraints. The large model is responsible for handling high-dimensional aesthetic and narrative topology planning, while the determined physical solver and database are responsible for taking over low-dimensional geometric anti-collision and connection tests. This mechanism of decoupling and checks and balances makes the system not only have human-like aesthetic ability, but also ensure 100% bottom-level geometric rationality.

4.5. Algorithm horizontal evaluation form

Based on the underlying technology evolution logic sorted out in this review, this paper summarises the mainstream algorithm of non-combat game building space generation as follows (Table 1):

Table 1. Comparison of the horizontal mechanism of the core algorithm of modern game building space generation

Driving mechanism	Algorithm	Aesthetic Fit	Connectivity	Technical bottleneck
Deterministic Rules	BSP tree, cellular automata, WFC	Extremely low	higher	It is prone to getting stuck in a "dead end" under high degrees of 3D constraints; it is unable to comprehend the high-level architectural semantics.

Table 1. (continued)

Feature Mapping	CNN, GAN, Diffusion model	polar altitude	Extremely low / Fragile	Lack of understanding of the macroscopic graph theory path; extreme reliance on massive, high-quality training data.
Deep RL / PCGRL	PPO Algorithm	secondary	polar altitude	3D spatial state explosion in dimension; the design of Reward Shaping is extremely difficult.
Neuro-Symbolic	LLMs	polar altitude	polar altitude	API calls and context window costs are high; there is still the occasional "spatial illusion" in large models.

From the above comparison table, it can be seen that the generation of game architectural space is going through an integration path from "fighting alone" to "take on strengths and make up for shortcomings". In the future, any industrial-level PCG engine cannot be supported by a single model alone. It must be a complex ecosystem in which multiple underlying drive mechanisms are constant and couple with each other.

5. Conclusions

This article systematically sorts out the evolution of data-driven game-building space generation technology. From the early rigid geometric rules to the feature mapping of the generative model, to the graph theory topology optimisation of RL, AI has shown the ability to realise the automatic generation of medium-complexity non-combat game scenarios. However, at the forefront of cross-research on game design and architecture, the current PCG engine still faces the following deep challenges that need to be broken through:

The explosion of the three-dimensional spatial state and the path of dimension reduction. At present, most mainstream RL solutions are limited to two-dimensional top-viewing planes or simple body-based grids. Considering the Z axis (multi-layered pavilion), the state space of MDP will grow rapidly to an unimaginable extent, and it is difficult to achieve instant computing on the current computer. The future research trend is to combine the dimension-down planning ability of LLMs with low-level physical solvers to guide the 3D physical instantiation of the bottom layer through high-level topological logic diagrams, so as to avoid the bottleneck of large-scale search directly in 3D state space.

Introducing deep semantic understanding and common sense reasoning in architecture. Although the Reward Shaping mechanism alleviates the problem of semantic missing to a certain extent, the existing Loss Function essentially still focuses on the purely mathematical relationship between "graphic connectivity" and physical geometry. Architectural design is not only the arrangement of space, but also a comprehensive carrier of human living habits, psychology and sociology. The breakthrough point in the future lies in the introduction of the common-sense reasoning ability of LLMs. Using the cross-modal semantic alignment of LLM, AI is expected to understand and solve multi-modal constraints with high subjectivity (such as using light topology to guide the player's vision) and realise the generation of high-rise buildings like real people.

From the delivery of physical functions to the generation of spatial narratives. The current AI generation technology is mainly in the basic stage of satisfying "aesthetic distribution fit" and "topological connection rate", which can only ensure the delivery of basic physical functions in space. From the dual perspective of architecture and game design, truly excellent spatial design requires the ability of "space specificity" and "Environmental Storytelling", that is, to create an "Internet celebrity punching point" with visual and emotional memory points.

However, the existing machine learning optimisation mechanism often rejects this. The mainstream loss function tends to erase the "inefficient space" or special nodes specially designed to render the plot atmosphere or assist the story scene. For example, the long and resourceless corridors designed to reflect the sense of oppression of the ruins are often eliminated as noise or punishment in the eyes of AI, which takes the efficiency of finding the way as the core indicator, causing the generated scene to fall into the homogenisation of mediocrity. As explored in the graphic narrative engine, future research should focus on how to quantify architectural semantics with emotional guidance and narrative attributes into neural network constraints that can be understood by AI. When AI can balance physical function and narrative specificity, programmed generation technology really has the soul of an architect.

References

[1] Bıyık, E. A., & Sürer, E. (2020). Developing a space syntax-based evaluation method for procedurally generated game levels. Mugla Journal of Science and Technology, 6(2), 1-10.
[2] Lech, B., Azad, S., Welnitz, J., Jonasson, J., & Martens, C. (2022). Designing a combined world and story procedural content generation engine. CEUR Workshop Proceedings.
[3] Snodgrass, S., Khalifa, A., Risi, S., Yannakakis, G. N., Liu, J., & Togelius, J. (2020). Deep learning for procedural content generation. Neural Computing and Applications, 33, 19-37.
[4] Earle, S., Edwards, A., Ahmed, I., & Togelius, J. (2023). Using unconditional diffusion models in level generation for Super Mario Bros. arXiv preprint arXiv: 2308.14247.
[5] Dai, S., Zhu, X., Li, N., Dai, T., & Wang, Z. (2024). Procedural level generation with diffusion models from a single example. Proceedings of the AAAI Conference on Artificial Intelligence, 38(16).
[6] Khalifa, A., Bontrager, P., Earle, S., & Togelius, J. (2020). PCGRL: Procedural content generation via reinforcement learning. Proceedings of the Sixteenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment.
[7] Kerker, L., Ingram, B., & Ranchod, P. (2025). Procedural generation of semantically correct levels in video games using reward shaping. School of Computer Science and Applied Mathematics, University of the Witwatersrand.
[8] Xu, K., & Verbrugge, C. (2025). A database-driven framework for 3D level generation with LLMs. arXiv preprint arXiv: 2508.18533.
[9] Todd, S., Earle, S., Nasir, M. U., Green, M. C., & Togelius, J. (2025). Video game level design as a multi-agent reinforcement learning problem. arXiv preprint arXiv: 2510.04862.
[10] Maleki, M. F., & Zhao, R. (2024). Procedural content generation in games: A survey with insights on emerging LLM integration. Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, 20(1), 167-178.
[11] Summerville, A., Snodgrass, S., Guzdial, M., Holmgård, C., Hoover, A. K., Isaksen, A., ... & Togelius, J. (2018). Procedural content generation via machine learning (PCGML). IEEE Transactions on Games, 10(3), 257-270.