

# PlotMap: Automated Layout Design for Building Game Worlds

Yi Wang, Jieliang Luo, Adam Gaier, Evan Atherton, Hilmar Koch  
Autodesk Research

{yi.wang, rodger.luo, adam.gaier, evan.atherton, hilmar.koch}@autodesk.com

**Abstract**—World-building, the process of developing both the narrative and physical world of a game, plays a vital role in the game’s experience. Critically-acclaimed independent and AAA video games are praised for strong world-building, with game maps that masterfully intertwine with and elevate the narrative, captivating players and leaving a lasting impression. However, designing game maps that support a desired narrative is challenging, as it requires satisfying complex constraints from various considerations. Most existing map generation methods focus on considerations about gameplay mechanics or map topography, while the need to support the story is typically neglected. As a result, extensive manual adjustment is still required to design a game world that facilitates particular stories. In this work, we approach this problem by introducing an extra layer of *plot facility layout design* that is independent of the underlying map generation method in a world-building pipeline.

Concretely, we define (*plot*) *facility layout tasks* as the tasks of assigning concrete locations on a game map to abstract locations mentioned in a given story (plot facilities), following spatial constraints derived from the story. We present two methods for solving these tasks automatically: an evolutionary computation based approach through Covariance Matrix Adaptation Evolution Strategy (CMA-ES), and a Reinforcement Learning (RL) based approach. We develop a method of generating datasets of facility layout tasks, create a gym-like environment for experimenting with and evaluating different methods, and further analyze the two methods with comprehensive experiments, aiming to provide insights for solving facility layout tasks.

**Index Terms**—World Building, Procedural Content Generation, Game Narrative, Reinforcement Learning, CMA-ES

## I. INTRODUCTION

Landscapes in digital games serve as more than just scenic backdrops; they interact intimately with the unfolding narrative, defining and shaping the player’s experience. Arming designers with tools considering narrative in map design, allows them to create more cohesive and immersive games. However, designing game maps is difficult, as it requires designers to consider varied qualities such as realistic topography [1], [2] and game playability [3] at the same time. Designing a map that supports a given story adds more constraints, making the problem even more challenging.

While the need to support an underlying story is typically neglected in most existing map generation methods, some efforts have been made to develop the story first, then to codify it as plot points and graphs so that maps can be generated based on their relations [4], [5]. However, as Dormans and Bakkes [6] pointed out, the principles that govern the design

of the spatial and the narrative side of the game are different, and thus these two processes should be independent. Methods for generating game maps from stories can appear artificially contrived to fit a story, and it is not straightforward to combine these methods with those that also take into account game design and geographical considerations. As a result, designing a game world that facilitates a story requires extensive manual modification; and as the number of constraints scale, the challenge of designing a map that satisfies all of the constraints of the story can become intractable, if not impossible, for a designer to do by hand [7].

We approach this problem by introducing an extra layer of *plot facility layout design* that is independent to the underlying map generation method in a world-building pipeline.

Our work is inspired by the philosophy behind the work from Dormans and Bakkes [6], which distinguishes the abstract space defined by the story (referred to as missions) and the concrete space defined by the actual geometric layout of a game map. The story is accommodated by mapping the former into the latter. While Dormans and Bakkes [6] focus on action adventure games with discrete “scenes” connected by “passages”, we impose very little assumption on methods used for story and map generation, and in particular target workflows for modern open world games.

We introduce the concept of *plot facilities*, which are abstract locations mentioned in the given story. A set of constraints are derived from the story in terms of the spatial relationships between plot facilities and elements in the underlying map. Given an underlying map, we arrange the layout of the plot facilities on top of the map to satisfy the constraints. Our method is compatible with any map generation technique in the sense that the underlying map can be hand-crafted, procedurally generated, or even from a Geographic Information System (GIS) such as Google Maps.

Figure 1 illustrates the world building approach and plot facility design task that is the focus of the work. We demonstrate our technique with a concrete pipeline using a procedural map generator and story constraints extracted from a textual story description with a large language model.

We consider two different representations of the underlying map: 1) a set of 2D polygons with precise geometric specifications, and 2) a pixel-based 2D image. We present two approaches to automatic plot facility layout design: an evolutionary computation (EC) based approach through Covariance Matrix Adaptation Evolution Strategy (CMA-ES)

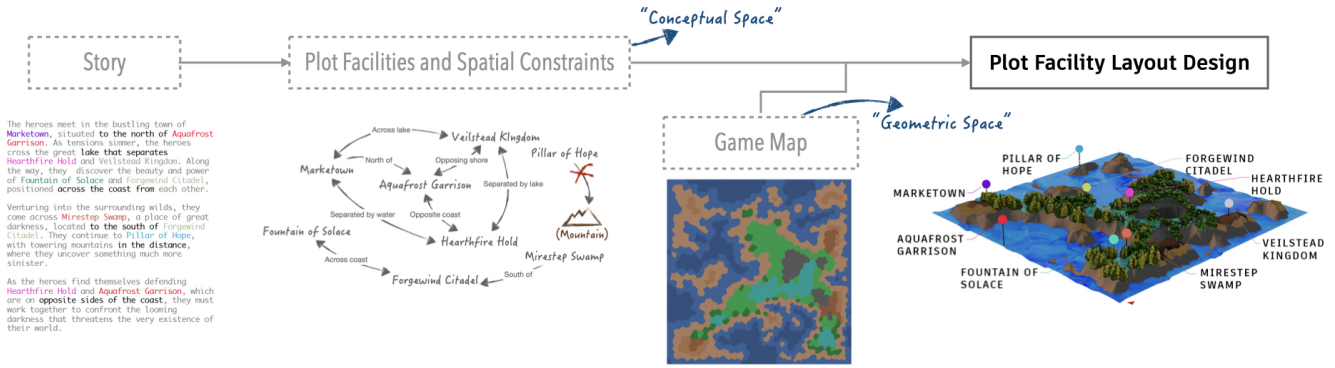


Fig. 1. Accommodating a story on a map with a plot facility layout design process

for a polygon-based map representation and a Reinforcement Learning (RL) based approach for maps as 2D images. We present baseline results for the two approaches. With extensive quantitative and qualitative studies, we show that the EC based approach can rapidly provide accurate solutions for tasks with a practical scale, while the RL based approach has the potential to better support human-in-the-loop map design workflows.

The paper also presents a dataset of facility layout tasks and a Gym-like environment to evaluate different methods and train the RL models. The dataset contains 10,000 plot facility layout tasks of different scales, involving 12 types of spatial constraints and maximum 60 plot facilities each task, based on a set of procedurally generated terrain maps.<sup>1</sup>

In summary, the paper’s contribution is threefold:

- We propose plot facility layout design as a novel approach to address the problem of supporting stories with game maps, which is compatible with most story and map generation methods.
- We provide a dataset of plot facility layout tasks and a Gym-like environment to experiment with and evaluate different methods for solving the task.
- We provide baseline results on a CMA-ES based method, and an RL approach for solving plot facility layout tasks, with discussions on their strengths and weaknesses.

## II. RELATED WORK

*a) Story and Game Map Generation:* Though intertwined, the generation of stories and maps are typically investigated in isolation. A few notable exceptions do tackle them as a single system. Hartsook et al. [5] proposed a story-driven map generation method where each event in the plot is associated with a location of a certain environment (e.g., castle, forest, etc.) and a linear plot line is translated to a constraint composed of a sequence of environments. Map generation is formulated as an optimization problem finding a topological structure of the map balancing requirements from a realistic game world and player preferences, subject to plot constraints. Valls-Vargas et al. [4] present a procedural method that generates a story and a map facilitating the story at the

same time. The problem is formulated as optimization of the topological structure of the map for both playability, and the space of possible stories.

Both Hartsook et al. [5] and Valls-Vargas et al. [4] generate grid-based maps consisting of discrete “scenes” connected by “passages”. The map structure is widely used in many classic games such as ROGUE [8] and early ZELDA series [9]. However, many modern RPG games feature seamless world maps with continuous terrains and very few geographical barriers for an immersive open world experience, such as ELDEN RING [10] and THE LEGEND OF ZELDA: BREATH OF THE WILD [11]. Dormans and Bakkes [6] use generative grammar based methods for both story (mission) generation and map (space) generation. The story elements are then mapped to spatial elements using heuristics specific to game genre. Our work establishes a mapping between narrative and spatial elements, but with a more general constraint satisfaction process.

*b) Procedural Content Generation:* Procedural Content Generation (PCG) has become an essential component in video games, employed for the algorithmic creation of game elements such as levels, quests, and characters. The primary objectives of PCG are to enhance the replayability of games, reduce the burden on authors, minimize storage requirements, and achieve specific aesthetics [12], [1], [2], [3]. Game developers and researchers alike utilize methods from machine learning, optimization, and constraint-solving to address PCG problems [13]. The primary aim of this work is to develop methods for plot facility layout design capable of generalizing across a wide range of environments and constraints. To achieve this goal, we employ a PCG approach to generate a diverse set of maps and constraints as our dataset.

*c) Evolutionary Computation for PCG:* Evolutionary computation (EC) reimagines game design challenges as optimization problems solved by evolving solutions to fulfill specific criteria [14]. The flexibility of EC allows for the creative redefinition of design challenges, unbounded by the constraints of precise problem formulations – for example, by redefining Mario levels as musical compositions and evolved neural networks suggest strategic tile placements [15].

In this work, we employ the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [16] for solving the plot facility layout task. CMA-ES is an algorithm designed to solve

<sup>1</sup>Our code and dataset are publicly available at: <https://github.com/AutodeskAILab/PlotMap>.

complex, non-linear, and non-convex optimization problems, by evolving a set of solutions, through adjusting a multivariate normal distribution, defined by mean and covariance, based on the success of past solutions. This process involves generating new solutions, evaluating them, and refining the distribution to improve outcomes. CMA-ES dynamically adjusts its search strategy, optimizing the balance between broad exploration and detailed examination of promising areas. It is highlighted for its efficiency, minimal parameter tuning, robustness against local optima and noisy evaluations, and suitability for a wide range of applications. However, its computational and memory demands increase quadratically with problem dimensionality, posing challenges for high dimensional (>1000D) problems.

*d) RL in Modern Video Games:* The popularity of games in AI research is largely attributed to their usefulness in the development and benchmarking of Reinforcement Learning (RL) algorithms [17]–[20]. On the other hand, RL can also be used as a design tool for modern games, especially for accessing and testing games. Iskander et al. [21] developed an RL system to play in an online multiplayer game alongside human players. The historical actions from the RL agent can contribute valuable insights into game balance, such as highlighting infrequently utilized combinations of actions within the game’s mechanics. Bergdahl et al. [22] used RL as an augmenting automated tool to test game exploits and logical bugs. Chen et al. [23] released a multi-agent RL environment to study collective intelligence within the context of real-time strategy game dynamics. To the best of our knowledge, this is the first instance that uses a learning-based approach to accommodate stories on game maps.

### III. PROBLEM FORMULATION

Put simply, we would like to assign every location mentioned in a story to an appropriate location on the game map (Fig. 2). Inspired by Dormans and Bakkes [6], we view the story and the geometric layout of a game map as independent of each other, except that the geometric layout should accommodate the story. We introduce the notion of *plot facilities*: conceptual “locations” mentioned in the story. These “locations” are *abstract* in the sense that they don’t correspond to any concrete geometric locations (yet). For example, the event “the hero finds an injured dwarf in the forest” happens at some place. There can be multiple locations on the map where this “abstract location” can be “instantiated”, as long as it does not contradict with the story - in this example it should be inside a forest area.

A set of constraints can be derived from the story for determining whether a particular layout of plot facilities is valid. The set of all plot facilities and the constraints form a conceptual space defined by the story, which is at a higher abstraction level than a concrete game map. The problem is then to assign geometric locations on the map to the plot facilities such that the constraints are satisfied - we call this problem *plot facility layout design*. A plot facility layout is essentially a mapping between the conceptual space defined by the story and the geometric space defined by the map.

In the following subsections, we describe our methods for this problem. To demonstrate a concrete map generation pipeline, in this study we specifically work with terrain maps, where different regions on the map represent different biomes.

#### A. Plot Facility Layout Design

We define a (*facility layout*) task as a tuple

$$\langle \mathcal{F}, \mathcal{T}, \mathcal{C} \rangle$$

where

- $\mathcal{F}$  is the set of (plot) facilities. Each facility has an identifier to be referred to by the constraints.
- $\mathcal{T}$  is an underlying terrain map.
- $\mathcal{C}$  is a set of spatial constraints over facilities in  $\mathcal{F}$  and biome types.

An (optimal) solution to a task is an assignment of coordinates to all the facilities in  $\mathcal{F}$ , so that a maximum number of the constraints in  $\mathcal{C}$  are satisfied considering their relations with each other and the biomes in  $\mathcal{T}$ .

In the subsequent sections, we consider two different representations of the map  $\mathcal{T}$ : 1) as a set of 2D polygons on the map, each associated with a biome type (e.g., OCEAN, PLAINS, etc.), or 2) as a 2D image, with color-coded regions representing biomes. We also consider two different representations of the constraints in  $\mathcal{C}$ : 1) as relational expressions such as `AcrossBiomeFrom(OCEAN, fordlen_bay, snapfoot_forest)`, or 2) as natural language (NL) utterances, such as “Fordlen Bay and Snapfoot Forest are separated by the ocean.”.

### IV. TASK DATASET GENERATION

We generate a dataset of 10,000 facility layout tasks for training and evaluation. Each task requires arranging the layout of minimum 10 to maximum 60 plot facilities on top of a procedurally generated map, w.r.t. a set of maximum 90 spatial constraints. Maps consist of 9 biome types and the constraints are generated based on 12 constraint types.

#### A. Map Generation

We employ a procedural map generation approach adapted from Patel [24] and its implementation by Dieu et al. [25]. Initially, a grid of Voronoi polygons is created from random points, followed by Lloyd relaxation to ensure even distribution. Coastline generation uses a flooding simulation to designate ocean, coast, lake, and land tiles based on water edge proximity and neighboring tiles. Elevation is determined by distance from the coast, with a distribution that emphasizes smoother terrain with fewer high points. Rivers flow from mountains to the nearest water body, while moisture levels are influenced by proximity to freshwater sources. Finally, the biome of each polygon is assigned based on a combination of moisture and elevation.

To generate a dataset, 100 maps with 1000 polygons each are produced. These maps are also converted to RGB images, suitable for input to neural network-based RL agents.




Fig. 2. We derive spatial constraints from a story and layout locations mentioned in the story on a map to satisfy the constraints.

### B. Constraint Generation

Facility layout tasks are generated by associating a set of random constraints to randomly chosen map. Figure 3 lists the constraint types and their frequency in the dataset. The constraint types were selected to cover common geometric relations between points (facilities) and polygons (biomes), and they are represented as unary, binary and ternary relations.

For each constraint type, we define a heuristic function for evaluating an existing facility layout w.r.t. any instantiation of the constraint type. The function returns a real number in  $[0.0, 1.0]$  with 1.0 meaning fully satisfied and 0.0 completely unsatisfied.<sup>2</sup> These functions are used to check if the randomly generated constraints are satisfied by a random layout and to compute the objective value or reward. We associate with each constraint a natural language utterance (e.g., “Fordlen Bay and Snapfoot Forest are separated by the ocean.”) for training the RL model. Tasks are then generated following Algorithm 1. Note that, as the constraints are extracted from an example layout, they are guaranteed to be solvable.

#### ALGORITHM 1: Facility Layout Task Generation

- Input:** A set of map  $MAP$ , maximum number of facilities  $N$ , a set of constraint types  $CT$ , minimum and maximum number of constraints  $M_1$  and  $M_2$
- Output:** A facility layout task  $\langle F, T, C \rangle$
1. Randomly sample a map  $T$  from  $MAP$ ;
  2. Randomly assign a location to facilities  $obj_1 \dots obj_N$  on the map  $T$ ;
  3. For each constraint type in  $CT$ , generate all possible instantiations of it w.r.t.  $obj_1 \dots obj_N$  and biome types. Evaluate each of them against the current map, adding the true ones to a set  $C'$ ;
  4. Sample a set of statements from  $C'$  sized between  $M_1$  and  $M_2$ , obtaining  $C$ ;
  5. For each statement in  $C$ , use large language model such as GPT [26] to rephrase it with a natural language sentence, resulting in a set of NL utterances  $C^{NL}$ ;
  6. **return** task  $\{\{obj_1 \dots obj_N\}, T, C^{NL}\}$ .

### V. EVOLUTIONARY COMPUTATION (CMA-ES) APPROACH

When the map is represented as polygons and the constraints as relational expressions, the facility layout problem can be approached as a black-box optimization problem where the goal

<sup>2</sup>E.g.,  $closeTo(x,y)$  is negatively correlated to the distance between  $x$  and  $y$ , and reaches 1 when their distance is less than a certain threshold.

Constraint Types	Meaning	Frequency
$AcrossBiomeFrom(b_1, p_1, p_2)$	$p_1$ is across biome $b_1$ from $p_2$	24895
$Outside(b_1, p_1)$	$p_1$ is outside biome $b_1$	5090
$Inside(b_1, p_1)$	$p_1$ is inside biome $b_1$	1117
$AwayFrom(b_1, p_1)$	$p_1$ is away from biome $b_1$	773
$CloseTo(b_1, p_1)$	$p_1$ is close to biome $b_1$	2362
$ToTheSouthOf(b_1, p_1)$	$p_1$ is to the south of biome $b_1$	739
$ToTheNorthOf(b_1, p_1)$	$p_1$ is to the north of biome $b_1$	746
$ToTheWestOf(b_1, p_1)$	$p_1$ is to the west of biome $b_1$	805
$ToTheEastOf(b_1, p_1)$	$p_1$ is to the east of biome $b_1$	727
$CloseTo(p_1, p_2)$	$p_1$ is close to $p_2$	103
$AwayFrom(p_1, p_2)$	$p_1$ is away from $p_2$	4215
$InBetween(p_1, p_2, p_3)$	$p_1$ is between $p_2$ and $p_3$	416
$OnSouth(p_1)$	$p_1$ is on the south of the map	283
$OnNorth(p_1)$	$p_1$ is on the north of the map	295
$OnEast(p_1)$	$p_1$ is on the east of the map	291
$OnWest(p_1)$	$p_1$ is on the west of the map	285
$ToTheSouthOf(p_1, p_2)$	$p_1$ is to the south of facility $p_2$	2774
$ToTheNorthOf(p_1, p_2)$	$p_1$ is to the north of facility $p_2$	2715
$ToTheWestOf(p_1, p_2)$	$p_1$ is to the west of facility $p_2$	2775
$ToTheEastOf(p_1, p_2)$	$p_1$ is to the east of facility $p_2$	2779
$VisibleFrom(p_1, p_2)$	$p_2$ is visible from $p_1$	616

Fig. 3. Constraint types included in the 10,000-task dataset (each  $p_i$  represents a plot facility). A constraint type  $ConstraintType(b_1, \dots, b_m, p_1, \dots, p_n)$  is instantiated to become a constraint by substituting each of  $b_1, \dots, b_m$  with a biome type, and each of  $p_1, \dots, p_m$  with a plot facility id ( $m \geq 0, n \geq 0$ ).

is to find facility positions, represented by  $(x, y)$  coordinates, to maximize constraint satisfaction. Each coordinate of each facility is a single dimension of search (e.g., the location of 10 facilities is a 20 dimensional problem). The search is guided only by feedback from calls to an objective function that tests the satisfaction of the entire layout. We use a common variant of CMA-ES, the Increasing Population Size CMA-ES [27] which restarts with a population size that doubles after each restart, based on predefined stagnation criteria.

#### A. Quantitative Evaluation

We evaluate the performance of CMA-ES on task configurations with 10, 30, and 60 facilities, each under conditions of 30, 60, and 90 constraints. For each configuration, we tested 1000 examples generated, as described in Section IV, allowing us to explore the impact of an increasing number of constraints on the optimization efficiency of a fixed number of facilities.

The plots in Figure 4 show that CMA-ES performs extremely well in low-dimensional settings, rapidly finding

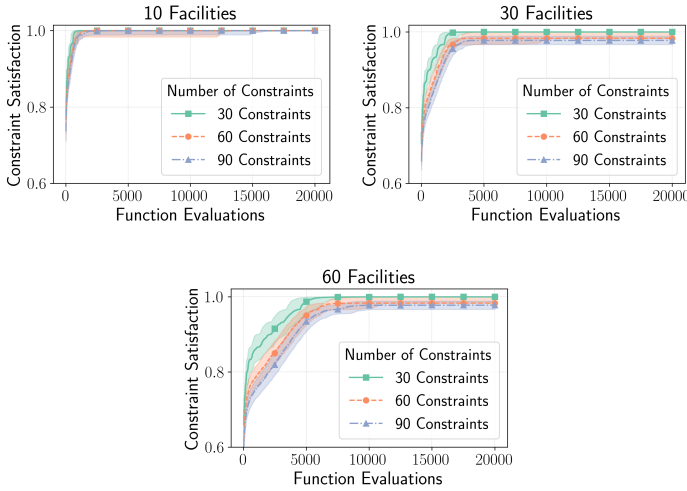


Fig. 4. Layout evaluations to reach a level of constraint satisfaction with 10, 30 and 60 facilities, and 30, 60, and 90 constraints to satisfy. Median satisfaction over 1000 different tuples of terrain and constraints, shaded regions indicate 1st and 3rd quartiles. Adaptive population sizes and cause generation times to vary from 0.25 to 0.75 seconds. However, performing a run with a budget of 20k evaluations consistently takes 2-5 minutes.

nearly perfect solutions even as the number of constraints grows. However, as the number of facilities grows, the number of evaluations required to reach a solution increases. The performance drop due to more constraints is low, highlighting the benefits of the black-box approach.

## VI. A REINFORCEMENT LEARNING APPROACH

Our CMA-ES based approach relies on repeated computation of the constraint satisfaction degree at solving time, which is enabled by the polygon-based representation of the maps. For many use cases, such a polygon-based representation of the map may not be available, for example when the map is hand-drawn or a photo of a physical spatial structure. Furthermore, the constraints may not be characterizable by closed-form arithmetic expressions. They may be from aesthetic considerations or subjective preferences that can only be statistically captured. For example, a game designer may want a “balanced” map layout. This motivates us to explore a learning-based approach for plot facility layout design.

In this section, we present a Reinforcement Learning (RL) method as a preliminary exploration. We employ a decision-making agent to optimize facility layout on a 2D pixel map image, considering the constraints of the story. This approach eliminates the need for solving-time constraint satisfaction computation, and allows for constraints to be expressed in natural language, enhancing the system’s flexibility and applicability to a wider range of design scenarios.

### A. RL Formulation

Each plot facility layout task can be viewed as sequentially moving a set of plot facilities  $\mathcal{F}$  on a map, thus we define the plot facility layout design as a sequential Markov decision process (MDP), with the following elements:

- Each state  $s \in S$  consists of three modalities: 1) a pixel-based image representing the map, 2) a real-valued vector

representing essential information of the plot facilities, and 3) a group of constraints.

- Each action  $a \in A$  is a 2d vector of real-valued  $[\Delta x, \Delta y]$  for one plot facility. In each round, plot facilities are moved one at a time in a fixed order, with the state and rewards updated after each movement. We set the range of  $\Delta x$  and  $\Delta y$  to be small steps so that we are simulating a concurrent movement of the facilities at a macro-level.
- The reward (for each step)  $r_t$  is +1 when all constraints are satisfied, and is the average satisfaction score from all constraints minus 1 when partial constraints are satisfied:

$$r_t = \begin{cases} 1, & \text{if all constraints are satisfied} \\ \frac{1}{n} \sum_{i=1}^n s_i - 1, & \text{otherwise} \end{cases}$$

where  $n$  is the number of constraints and  $s_i$  is the satisfaction score for each constraint. The satisfaction score for each type of constraint is within  $[0, 1]$  and defined based on the heuristic functions described in Section IV-B. The range of the reward  $r_t$  is  $[-1, 0] \cup \{1\}$ .

- The transition function is deterministic, where  $s_{t+1} = f(s_t, a_t)$ .
- Each episode is terminated when all the constraints are satisfied or at 200 timesteps.

We train an RL agent to learn an optimal policy  $\pi_\theta$  in order to maximize the expected discounted rewards:

$$\max_{\pi_\theta} \mathbb{E}_{\tau \sim \pi_\theta} \left[ \sum_{t=0}^T \gamma^t r(s_t, a_t) \right], \quad (1)$$

where trajectory  $\tau = (s_0, a_0, s_1, a_1, \dots, s_T, a_T)$ ,  $\theta$  is the parameterization of policy  $\pi$ , and  $\gamma$  is the discounted factor. Our goal is to train a general RL agent to be able to find solutions to any arbitrary task.

## VII. EXPERIMENTS & ANALYSIS

### A. Experiment Setup

Our state space includes 3 types of inputs:

- a map encoded as a  $(42 \times 42 \times 3)$  RGB image.
- constraints represented by natural language utterances or relational expressions, depending on the embedding strategies described in the following paragraph.
- Each plot facility’s information is represented by a vector, consisting of its position  $[x, y]$  on the map, a binary motion indicator signifying if it is its turn to move or not, and a unique identifier.

In all of our experiments, we have set a limit of 10 for both the number of plot facilities and the number of constraints. All of our policies are trained using Proximal Policy Optimization (PPO) [28]. To handle the multi-modal observation space, we employ pre-trained models to independently extract embeddings from the map (image) and constraint (natural language) inputs. These embeddings are subsequently concatenated with the informational vector and provided as inputs to the policy network. Specifically, we design two strategies for deriving the embeddings:

- *NL-based*: using ResNet [29] for maps and Sentence-Transformer [30] for the NL representation of constraints. The state dimensions are 4,392, consisting of 512 dimensions for map, 40 for plot facilities, and  $384 \times 10$  for constraints.
- *Relation-based*: using ResNet for maps, and each constraint, as a relational expression, is encoded as a one-hot vector, representing the constraint type, followed by three one-hot vectors indicating the specific plot facilities to instantiate the constraint with. The state dimensions are 1,782, consisting of 512 dimensions for map, 40 for plot facilities, and  $123 \times 10$  for constraints.

### B. Quantitative Evaluation

We examine the performance of our two proposed embedding strategies across four small task sets. For the datasets comprising 100 tasks, we explore the influence of maps and constraints on generalization. These results are presented in Table I.

For each dataset, we evaluate our trained policies under two conditions: *rand init*, which refers to evaluating the policies with the same task sets used for training but under different initial positions; and *rand init + unseen tasks*, which reports the success rates when the policies, each trained on their respective task set, are tested on 100 unseen tasks (with different maps and constraints). Success is considered only when all constraints are satisfied. We calculate the success rates over 1,000 rollouts, with each rollout taking approximately 5 seconds to complete.

Table I shows that both baselines perform exceptionally well on tackling a single hard task. The success rate for both baselines decreases as the size of the task set increases, exhibiting a deficiency in generalization capability.

To investigate further the factors impeding generalization, we study the influence of maps and constraints by training on three distinct sets of 100 tasks as depicted in the last three columns in Table I. The substantial disparity between varying sets of constraints and one fixed set of constraints indicates that constraints pose a greater challenge to generalization, regardless of the encoding method employed.

## VIII. QUALITATIVE RESULTS

In this section, we demonstrate a complete pipeline of the system, and highlight insightful behaviors of the CMA-ES and RL method through specific examples.

Figure 2 illustrates a complete story-to-facility-layout pipeline. From the story on the left, 8 plot facilities and 6 constraints are extracted with a pre-trained large language model.

The map on the right shows the resulting plot facility layout, along with motion trails indicating traces to the final locations from rolling out a trained RL model. The RL models demonstrate interesting cooperative behaviors (Fig. 5).

In Figure 6, we demonstrate the same story accommodated on 4 different maps. We observe that the same set of plot facilities can still maintain their relative spatial relations on




Fig. 5. **Cooperative behavior to satisfy constraints**: Markettown and Veilstead Kingdom must be across a lake from each other, while Aquafrost Garrison must be to the south of Markettown. Aquafrost Garrison was initially to the south of Markettown but as Markettown moves south to be across the lake from Veilstead Kingdom, Aquafrost Garrison moves even further south to continue satisfying its south of Markettown constraint.

completely different geometric layouts, which aligns with the perspective described in [6]: “The same mission (story) can be mapped to many different spaces, and one space can support multiple different missions (stories)”. This capability enables designers to envision various interpretations of unspecified story details and potential story progressions.

The RL method provides smooth traces from the starting positions to the final layout, enabling an interactive interface for adjusting existing layout solutions. In Figure 7 (a), we show that our RL policies support real-time re-adaptation after human intervention. Specifically, after we manually change the location of Markettown to the northeast part of the map, Veilstead Kingdom and Hearthfire Hold can adjust their locations to continue to be across the lake from Markettown, while Aquafrost Garrison stays at the same location, so all of the constraints are still satisfied. In contrast, Figure 7 (b) shows the the result obtained by CMA-ES, again with a manual change of the location of Markettown. For the rest of the facilities, we start with their current location as the initial guess, and re-run the CME-ES process until reaching 1.0 satisfaction. As can be seen, the configuration of the rest of the facilities has been drastically changed, with no obvious connection with the previous layout.

## IX. LIMITATION AND FUTURE WORK

The scalability of CMA-ES is limited by the dimensionality of the problem and the complexity of the objective function. The quadratic scaling of the size of the covariance matrix dramatically reduces the speed of CMA-ES with a large number of parameters, and CMA-ES is not often used in cases where it exceeds 1000. This translates to 500 locations as each location is represented by two parameters (x,y). The 2011 release SKYRIM had  $670^3$ , the more recent ELDEN RING had a map twice the size with locations at higher density<sup>4</sup>. In the

<sup>3</sup><https://screenrant.com/skyrim-how-many-locations-big-map-size/>

<sup>4</sup><https://screenrant.com/skyrim-map-elden-ring-lands-between-big-size>

TABLE I

SUCCESS RATE (%) COMPARISON BETWEEN A RANDOM AGENT AND THE TWO PROPOSED BASELINE METHODS ACROSS TASK SIZES AND CONDITIONS. EACH BASELINE IS ASSESSED UNDER TWO CONDITIONS: 1) RANDOM INITIALIZATION ON THE SAME TRAINING TASKS, AND 2) RANDOM INITIALIZATION ON 100 UNSEEN TASKS. FOR THE 100-TASK DATASETS, WE EXAMINE THREE COMBINATIONS OF MAPS AND CONSTRAINTS.

Method	Evaluation Condition	Success Rate (%)					
		1 task	5 tasks	50 tasks	10 maps + 100 sets of constrs	100 tasks 1 map + 100 sets of constrs	100 maps + 1 set of constrs
Random Agent		23.9	30.3	30.5	30.9	34.5	38.0
NL-based	rand init	84.1	48.2	38.5	38.4	42.8	79.2
	rand init + unseen tasks	34.8	33.3	35.1	38.0	39.9	70.4
Relation-based	rand init	100.0	55.1	46.1	38.5	48.1	77.5
	rand init + unseen tasks	32.2	29.6	32.7	36.1	40.1	68.9




Fig. 6. Same story accommodated on 3 different maps

truly large scale cases where CMA-ES would be most useful it is also the least practical.

Our dataset generation procedure resulted in an imbalanced dataset, as shown in Figure 3, which has led to unexpected behaviors and failure cases of the trained RL models. For example, in the dataset there is a significantly higher proportion of AcrossBiomeFrom constraint. This had led to a strong tendency of the RL model to move facilities to the edge of the map, since two plot facilities on different edges of the map are likely to be across several different biomes from each other, despite human designers may prefer the facilities to be close to the biome.<sup>5</sup> The RL model also tends to fail to satisfy constraint types with fewer occurrences in the dataset (such as CloseTo( $p_1, p_2$ )). The behavior of the CMA-ES method is not affected by the imbalance of the dataset. However, it could still produce undesirable maps due to the inconsistency between the human designer’s interpretation and the computational implementation of a constraint.

The use of hand-crafted objective functions present challenges in both accurately reflecting the preferences of human designers and providing the right signal for training. A promising solution for better understanding the kinds of solutions which designers prefer is to adopt a statistically trained preference model such as used in Reinforcement Learning from

<sup>5</sup>For example, “A and B are across a lake” generally implies that A and B are on the shore of the lake.

Human Preference, which offers potentially more accurate reflections of human intent [31]. Moreover, we considered the satisfaction of all constraints as the benchmark for successful task resolution. In practice, a suboptimal solution that satisfies most of the constraints might be acceptable, and situations where the constraints are unsatisfiable are completely possible. In these cases knowledge of designer preferences, or a mixed-initiative approach [32]–[34] which allowed editing of the map, would allow desirable solutions to be found.

The formulation, scalability, and embedding strategies of the RL approach all presented challenges. Employing a single RL agent to manage all global information and plot facilities is a bottleneck, hindering performance in larger settings. A distributed RL model, where each facility operates as an independent agent could allow the system to better scale. Achieving a high level of generality within the RL model, akin to adapting to various game levels or rules, poses a significant challenge. Drawing parallels with RL applications in gaming, where agents must learn to generalize to varied conditions, a curriculum learning approach [35] could offer a solution by dynamically generating maps and constraints while progressively increasing scenario complexity and diversity, thereby fostering the model’s generalization capabilities.

## X. CONCLUSION

In this work we introduced new methods to support stories with game maps through an automated plot facility layout

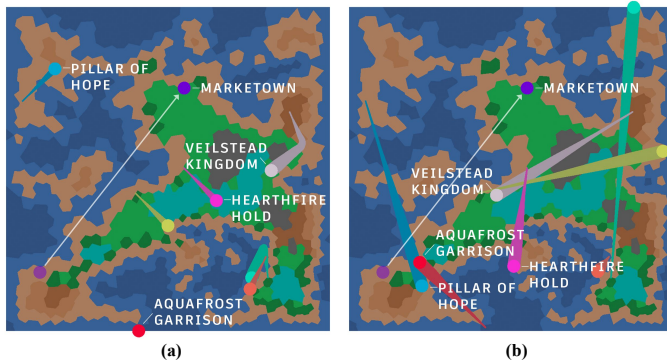


Fig. 7. Plot facility re-adaptation after the user moves Markettown to a different location (a) with RL, and (b) with CMA-ES. (a) can be seen as an adjustment of the previous layout while (b) changes into a completely different layout.

design process. We demonstrated that this approach allows us to utilize existing story and map generation techniques, and visualize the spatial implications of a story and the narrative potential of a map. We present two methods for solving plot facility layout tasks: an EC based approach through CMA-ES, and an RL based approach, suitable for different representations of the underlying maps. With extensive quantitative and qualitative studies, we show that the EC based approach can rapidly provide accurate solutions for tasks with a practical scale, while the RL based approach has the potential to better support human-in-the-loop map design workflows. The concept of plot facility layout design has potential in many game design applications, such as map design, playtime quest generation/adaptation and story debugging; but also potential applications in other domains involving spatial layouts subject to constraints, such as the design of large office buildings or manufacturing plants.

## REFERENCES

- [1] R. Smelik, K. J. de Kraker, S. Groenewegen, T. Tutenel, and R. Bidarra, "A Survey of Procedural Methods for Terrain Modelling," Jun. 2009.
- [2] G. Kelly and H. McCabe, "A Survey of Procedural Techniques for City Generation," *The ITB Journal*, vol. 7, no. 2, May 2017. [Online]. Available: <https://arrow.tudublin.ie/itbj/vol7/iss2/5>
- [3] R. van der Linden, R. Lopes, and R. Bidarra, "Procedural generation of dungeons," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 6, no. 1, pp. 78–89, 2014.
- [4] J. Valls-Vargas, S. Ontanón, and J. Zhu, "Towards story-based content generation: From plot-points to maps," in *2013 IEEE Conference on Computational Intelligence in Games (CIG)*. IEEE, 2013, pp. 1–8.
- [5] K. Hartsook, A. Zook, S. Das, and M. O. Riedl, "Toward supporting stories with procedurally generated game worlds," in *2011 IEEE Conference on Computational Intelligence and Games (CIG'11)*. IEEE, 2011, pp. 297–304.
- [6] J. Dormans and S. Bakkes, "Generating missions and spaces for adaptable play experiences," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 3, no. 3, pp. 216–228, 2011.
- [7] R. Matsumoto, "Introduction of case studies of engineer efforts to respond to the open field of elden ring," *Computer Entertainment Developers Conference*, 2022.
- [8] A. Design, "Rogue," 1980.
- [9] Nintendo, "The legend of zelda," 1986.
- [10] FromSoftware, "Elden ring," 2022.
- [11] Nintendo, "The legend of zelda: Breath of the wild," 2017.
- [12] M. Hendriks, S. Meijer, J. Van Der Velden, and A. Iosup, "Procedural content generation for games: A survey," *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 9, no. 1, pp. 1–22, 2013.

- [13] J. Togelius, G. N. Yannakakis, K. O. Stanley, and C. Browne, "Search-based procedural content generation: A taxonomy and survey," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 3, no. 3, pp. 172–186, 2011.
- [14] J. Liu, S. Snodgrass, A. Khalifa, S. Risi, G. N. Yannakakis, and J. Togelius, "Deep learning for procedural content generation," *Neural Computing and Applications*, vol. 33, no. 1, pp. 19–37, 2021.
- [15] A. K. Hoover, J. Togelius, and G. N. Yannakakis, "Composing video game levels with music metaphors through functional scaffolding," in *First computational creativity and games workshop. ACC*, 2015.
- [16] N. Hansen and A. Ostermeier, "Completely derandomized self-adaptation in evolution strategies," *Evolutionary computation*, vol. 9, no. 2, pp. 159–195, 2001.
- [17] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling, "The arcade learning environment: An evaluation platform for general agents," *Journal of Artificial Intelligence Research*, vol. 47, pp. 253–279, 2013.
- [18] M. Jaderberg, W. M. Czarnecki, I. Dunning, L. Marris, G. Lever, A. G. Castaneda, C. Beattie, N. C. Rabinowitz, A. S. Morcos, A. Ruderman *et al.*, "Human-level performance in 3d multiplayer games with population-based reinforcement learning," *Science*, vol. 364, no. 6443, pp. 859–865, 2019.
- [19] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev *et al.*, "Grandmaster level in starcraft ii using multi-agent reinforcement learning," *Nature*, vol. 575, no. 7782, pp. 350–354, 2019.
- [20] C. Berner, G. Brockman, B. Chan, V. Cheung, P. Debiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse *et al.*, "Dota 2 with large scale deep reinforcement learning," *arXiv preprint arXiv:1912.06680*, 2019.
- [21] N. Iskander, A. Simoni, E. Alonso, and M. Peter, "Reinforcement learning agents for ubisoft's roller champions," *arXiv preprint arXiv:2012.06031*, 2020.
- [22] J. Bergdahl, C. Gordillo, K. Tollmar, and L. Gisslén, "Augmenting automated game testing with deep reinforcement learning," in *2020 IEEE Conference on Games (CoG)*. IEEE, 2020, pp. 600–603.
- [23] H. Chen, S. Tao, J. Chen, W. Shen, X. Li, S. Cheng, X. Zhu, and X. Li, "Emergent collective intelligence from massive-agent cooperation and competition," *arXiv preprint arXiv:2301.01609*, 2023.
- [24] A. Patel, "Polygonal map generation for games," *Red Blob Games*, vol. 4, 2010.
- [25] D. Dieu, M. Markiewicz, K. Grodzicki, and SWi98, "Polygonal map generation for games," <https://github.com/TheFebriin/Polygonal-Map-Generation-for-Games>, accessed 2023.
- [26] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [27] A. Auger and N. Hansen, "A restart cma evolution strategy with increasing population size," in *2005 IEEE congress on evolutionary computation*, vol. 2. IEEE, 2005, pp. 1769–1776.
- [28] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [29] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [30] N. Reimers and I. Gurevych, "Sentence-bert: Sentence embeddings using siamese bert-networks," *arXiv preprint arXiv:1908.10084*, 2019.
- [31] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, "Deep reinforcement learning from human preferences," *Advances in neural information processing systems*, vol. 30, 2017.
- [32] G. N. Yannakakis, A. Liapis, and C. Alexopoulos, "Mixed-initiative co-creativity," in *International Conference on Foundations of Digital Games*, 2014.
- [33] G. Smith, J. Whitehead, and M. Mateas, "Tanagra: A mixed-initiative level design tool," in *Proceedings of the Fifth International Conference on the Foundations of Digital Games*, 2010, pp. 209–216.
- [34] M. Cook, J. Gow, G. Smith, and S. Colton, "Danesh: Interactive tools for understanding procedural content generators," *IEEE Transactions on Games*, vol. 14, no. 3, pp. 329–338, 2021.
- [35] R. Portelas, C. Colas, L. Weng, K. Hofmann, and P.-y. Oudeyer, "Automatic curriculum learning for deep rl: A short survey," in *IJCAI 2020-International Joint Conference on Artificial Intelligence*, 2021.