

L-MCTS Framework of Massive Spatio-temporally Plausible Content Generation

Chia-Hsing Chiu, Yu-Chi Lai, Andrew Chen, Pai-Yi Su, Wei-Yao Ku, Wen-Kai Tai, and Shih-Syuan Lin

Abstract—Spatio-temporal content, cascades of physical and procedural events, is one of the most indispensable elements in games and animations nowadays. Generally, designing such scenes needs expertise to achieve both procedural rules and targeting criteria such as complexity and difficulty, thus being time-consuming and manpower intensive. Therefore, we propose a general framework that combines the L-system and the Monte Carlo Tree Search (MCTS), named L-MCTS, whose automatic workflow, integrating expertise and domain knowledge, is capable of massive production with targeted criteria. More specifically, we break the content space of near-infinity degrees of freedom into three parts: forming cascading events with L-system, positioning elements with inverse embedding, and activating the content with initial conditions. Furthermore, we adjust Monte Carlo Tree Search (MCTS) to systematically and coherently sample plausible instances through the decomposed space. Finally, we demonstrate its generality of our framework with three different applications and evaluate their effectiveness and efficiency through several experiments and user studies.

Index Terms—L-system layout representation, Monte Carlo Tree Search generation and evaluation, difficulty evaluation, procedural content generation, layout generation

I. INTRODUCTION

Spatio-temporal content driven by cascading events is commonly seen in games and animations. For examples, the chain reaction is such a visual treat that leaves people breathless when components trigger one another (left of Fig. 1), and the billiard trick shots that strike all object balls into the pockets in one shot is also appealing to the audience hence being referenced as "artistic pool" (right of Fig. 1). However, generating such content with physical plausibility is challenging, particularly when additional indicators such as appealingness or difficulties must be considered in practice. Furthermore, designing such content usually requires expertise and domain knowledge, and thus, it is difficult to automate the workflow. All the mentioned issues make the spatio-temporal content generation a manpower-intensive and time-consuming task in the fast-paced multimedia industry. Therefore, this work aims at developing a general framework that automatically generates massive amounts of plausible content without the presence of experts.

Content modeling and animation synthesis formulate space-time constraints as optimization problems, and Monte Carlo Optimization (MCO) is one of the successful solver for modeling trees and buildings [1], [2], generating levels and layouts [3]–[7], and synthesizing space-time animations [8]–[10]. However, random exploration of the plausible space is time-consuming. Therefore, our framework systematically




Fig. 1. This shows the examples of our targeting spatio-temporal content, where the left shows the chain reaction that propagates the interaction through various components and the right shows the "artistic pool" that challenges the player to strike all the objects balls in one shot.

parameterizes the spatio-temporal content in three aspects: to describe and deduce the element-wise interactions, to describe the element-wise spatial information with the geometric parameters, and to trigger the cascading events with possible initial conditions. In order to search for plausible spatio-temporal content, we also involve the simulation phase, which verifies validness and objectiveness of the content, resulting in a 4-level Monte Carlo Tree Search, named as L-MCTS. However, it is impractical and inefficient to search the extremely high-dimensional geometric parameters. Thus, we involve the domain knowledge to develop inverse embedding for narrowing parameter searching into a smaller space. Even though our L-MCTS allows an arbitrary function of plausibility, designing a proper metric to reflect objectiveness, such as appealingness, is still an unsolved problem. We free ourselves from domain dependency by using expertise expression and questionnaire postulation to reveal subjective opinions along with domain-free statistics, such as tree complexity and parameter variations.

After conceptual deduction and development, we have applied our framework on three applications, one-shot-all-eliminate breakout clone, chain reaction, and one-shot-all-sink billiard while having experiments and user studies to determine their proper parameters. The implementation can massively generate breakout levels, chain reaction setups, and billiard layouts with the target degree of the level difficulty, visual complexity, and playing difficulty, respectively. While comparing against other baselines, our framework can effectively enhance the search efficiency for content generation. We have also verified the effectiveness of our designed visual




Fig. 2. Our L-MCTS uses MCTS to hierarchically explore the spatio-temporal content space. The system is divided into 4 phases from left to right, **L-string expansion**, which describes and deduces the sequential interactions, **spatial positioning** that applies inverse embedding for better searching efficiency, **activating initialization** which picks up the initiation conditions, and **validation simulation** that evaluates the generated content, while the evaluation result is backpropagated to the tree for concentrating searching on more promising parameters.

complexity and difficulty metric through a set of user studies. Overall, our framework achieves the following contributions.

- 1) We propose **L-MCTS**, which incorporates L-system and Monte Carlo Tree Search (MCTS), as a general framework to automatically generate massive amounts of spatio-temporal content that fulfills the user-defined objectiveness.
- 2) We further enhance the search efficiency by decomposing the spatio-temporal content into the L-system-described cascading sequence, spatially varied geometric parameters, and activated initial conditions for systematic exploration.

After all, our system can effectively and automatically generate spatio-temporal instances under user specified degree of criteria.

II. RELATED WORKS

The proposed L-MCTS framework aims at automatic generation of various spatio-temporal content. Therefore, the following focuses on the core related works of content generation, which include L-system-based procedural modeling, space-time constraint simulation, and Monte Carlo optimization.

L-system-based procedural modeling first comes into our mind. Research [11]–[13] apply rewriting and turtle interpretation of L-system grammars to describe various branching plants and others [14], [15] extend the concept to describe variant objects. It is compact and able to construct various models with a small set of derived parameters along with its goodness at systematic and hierarchic deductions of all possible variations. However, to the best of our knowledge, there are only few grammar descriptions and production rules designed for temporal events [5]. Since spatio-temporal content consists of sequences of semantic interactions, for example, the sequence of ball-to-ball and ball-to-cushion collisions in billiard, it seems to be reasonable to derive L-system representations for these interactions because of its easy deduction. However, there are still challenges listed as the following. Traditional L-system only records interactions and components involved, but it does not consider their spatial positioning and temporal actions. Even though its variation, the parametric L-system, is capable of representing the continuous parameters through

the production rules, the complex physical phenomena makes plausible parameters and initial conditions seeking difficult and even impossible. Therefore, we adapt Monte Carlo Tree Search (MCTS) to explore the possible spatio-temporal space to seek reasonable solutions.

Space-time constraint simulation aims at achieving key-framed design and control using space-time constraints in order to generate artistically satisfied animations. There are various applications, including controlling multi-body trajectories [8]–[10], [16], [17], manipulating skeleton motions [18]–[20], editing deformable objects [21], [22], and controlling fluid [23], [24]. Generally speaking, the constraints are achieved by either local greedy control [16], [20], [24] or global optimization [18], [22], and we summarize them as the inverse solutions to model the content distribution with pre-defined constraints. However, these methods generally require domain knowledge to formulate domain-dependent solutions, i.e., they cannot be easily transferred to other applications as mentioned by Cheney *et al.* [8]. Different from the aforementioned methods, our system first designs forward sampling of cascading interactions, interaction-wise geometric parameters, and initial conditions as a general and domain-free framework. At the same time, we take inverse postulation to derive inverse embedding with simplified procedure models for interaction-wise prediction in order to accelerate the searching procedure. Furthermore, past studies target at generating one artist-satisfied animation, but ours is intended for massive production of animations achieving certain target goals with variation.

Monte Carlo Optimization (MCO), including Monte Carlo Tree Search (MCTS), obtains the optimal result with designated strategies using random sampling [25]. This concept is widely used in various domains such as the distribution of trees and buildings [1], [2], music composition [5], 2D platform level design [3], [6], collision layout synthesis [4] and the consensus threshold of preference relations [26]. However, randomly jumping around the plausible space is inefficient and cannot further explore those possibly important regions for traditional MCO. Therefore, Cheney *et al.* [8] apply Markov Chain Monte Carlo sampling in the parameter space for better efficiency while Anderson *et al.* [17] extend the concept

to flock simulation. MCTS is also introduced for coherent exploration and acceleration. The most famous application lies in the opponent AI of Go for AlphaGo and AlphaGo Zero [27]. In addition, Qi *et al.* [28] propose a fast MCTS pruning algorithm to achieve efficient spacecraft trajectory optimization, and Hong *et al.* [29] propose a sampling strategy for decision variables to solve the large-scale multiobjective optimization via MCTS. We summarize the above methods as the forward solutions to model the content distribution based on sampling strategies. As far as we know, there exists no work representing the spatio-temporal content for forward methods. Thus, we develop an L-system representation to encode the interactions and embed geometric and physical parameters to construct the possible spatio-temporal space in order for coherent and thorough exploration. Furthermore, we also enhance search efficiency by incorporating the inverse concept to have a hybrid solution, which have the benefits of both the forward generality and inverse efficiency by reducing the effectively searching domain.




Fig. 3. This shows two spatio-temporal examples of domino toppling and billiard shooting. The content generally consists of 4 parts, which are (a) **elements**, such as dominoes and balls, (b) **interactions**, such as domino topplings and ball collisions, (c) **geometric configurations**, such as the position of dominoes and balls, and (d) **initial conditions**, such as the initial force of toppling and striking.

III. PROBLEM FORMULATION

This work aims to design a fully-automatic framework which is capable of massive generation of spatio-temporal content given a targeted goal. Similar to space-time constraints [18], this work defines the spatio-temporal content as cascading interactions, denoted as **interactions**, between objects, referred as **elements**, where **elements** are placed on designated positions, denoted as **geometric configurations**, and the interactions follow designated procedural rules such as physics while triggering with designed **initial conditions** in the beginning. The degree of goal achievement, $f(\cdot)$, is evaluated based on the process of interactions from triggering until all objects stop moving, and it is fully application dependent while its evaluation criteria, which are usually empirical, requires domain knowledge and expertise. Furthermore, this work sets it to 0 when an interaction fails. Fig. 3 shows two examples, a domino run and a one-shot-all-sink billiard set-up. In a domino run, the process is defined as a success if all dominoes are toppled in once, or to be a failure vice versa. Similarly, a successful one-shot-all-sink billiard set-up is determined by whether all object balls are sunk into the pockets in one shot or not. Furthermore, a metric can be designed to

determine whether the set-up reaches the target goal or not, such as difficulty for the billiard set-up. To simplify the later discussion, we denote the spatio-temporal space, consisting of all plausible instances, \mathbf{a}^i , as $\mathcal{A} = \{\dots, \mathbf{a}^i, \dots\}$ whose $\mathbf{a}^i = \langle \mathbf{e}^i, \mathbf{t}^i, \mathbf{x}^i, \mathbf{z}^i \rangle$, where $\mathbf{e}^i, \mathbf{t}^i, \mathbf{x}^i, \mathbf{z}^i$ represent the instances of **elements**, **interactions**, **geometric configurations**, and **initial conditions**, respectively. Also, the evaluation metric is written as $f(\mathbf{a}^i)$ for any given instance \mathbf{a}^i , and the desired content is referred to that satisfies the equation $f(\mathbf{a}^i) > \lambda$, for λ being the targeted goal.

Accordingly, this work treats the targeted generation as sampling from the content space. To be specific, a plausible candidate could be any kind of setting under the application scope, including those that may provide undesired or even failure results. In the above exemplar domino run, a possible solution implies a set-up with any given domino number arranging in any designed pattern, even those fail to topple all dominoes in once or fail to meet the target standard. Since the involved number of objects, the placed locations, and etc. are all free parameters, the plausible content space is extremely high dimensional. Therefore, this makes it difficult to find solutions for a targeted goal without systematic exploration. As a result, the next section first describes the concept to semantically and parametrically decompose the content space for systematic and coherent exploration. Then, we adapt Monte Carlo Tree Search (MCTS) to efficiently sample the decomposed content space for efficient solution seeking. Finally, we apply the concept on three applications.

IV. PLAUSIBLE CONTENT SAMPLING

This work rephrases the targeted problem as a sampling problem which can be clearly stated as "how to sample a desired instance from \mathcal{A} in a massive amount". The preference of an instance is evaluated with $f(\mathbf{a}^i)$ for any given \mathbf{a}^i , while such a metric completely depends on the goal to achieve and is not restricted to a specific equation. For such a problem, we follow the approach proposed by Chenney *et al.* [8] and let $p_w(\mathbf{a}^i)$ model the probability that the content instance \mathbf{a}^i might arise in the simulation environment and $p_w(\mathbf{a}^i | f(\mathbf{a}^i) > \lambda)$ for those that even meet the user desire. We also follow their derivation to have

$$p(\mathbf{a}^i) \propto p_w(\mathbf{a}^i) p_f(\mathbf{a}^i) \quad (1)$$

where $p(\mathbf{a}^i)$ models the distribution of the plausible and desired instances in the content space, and $p_f(\mathbf{a}^i)$ depends only on how well the content satisfies objectiveness. According to the definition of \mathbf{a}^i and the independencies among its compositions in general, the equation can be further rewritten as

$$p(\mathbf{a}^i) \propto p_w(\mathbf{l}^i) p_w(\mathbf{x}^i) p_w(\mathbf{z}^i) p_f(\mathbf{a}^i). \quad (2)$$

where \mathbf{l}^i is the combination of \mathbf{e}^i and \mathbf{t}^i . This work proposes a new framework, named L-MCTS, which incorporates the L-system into our adapted Monte Carlo Tree Search (MCTS), to model $p(\mathbf{a}^i)$. Intuition behind this design is two folds. On one hand, experts usually design such content through a systematic mechanism to semantically compose relative parameters instead of sampling these parameters independently.

That is, both the **elements** and **interactions** are arranged into the chronological order while additional **elements** as well as **interactions** are inserted into the current sequence in an iterative fashion. Such a definition not only helps to sample e^i and t^i with unconstrained dimensionality but also better present the progressive process of the content design. On the other hand, modelling each component independently is pointless due to the dependencies between the components. For example, a domino could never topple its next target if it is placed too far away or not on the falling direction of the previous domino. As a result, the L-MCTS is designed to perform semantic and hierarchical exploration of the content by incorporating the L-system representation while combining inverse solutions during the search to greatly improves the search efficiency. More specifically, the search tree is separated into four parts: the **L-string expansion** that models $p_w(\mathbf{l}^i)$, **spatial positioning** that models $p_w(\mathbf{x}^i)$, **activating initialization** that models $p_w(\mathbf{z}^i)$ and **validation simulation** that models $p_f(\mathbf{a}^i)$. With the backpropagation mechanism, L-MCTS concentrates on more promising sub-trees and models $p(\mathbf{a}^i)$ efficiently. We show the algorithm of L-MCTS in Algorithm 1.

Algorithm 1 L-system Monte Carlo Tree Search (L-MCTS)

```

1:  $\mathcal{A}_{desired} = \{\}$ 
2: repeat
3:    $\mathbf{a}^i = \text{SelectToLeaf}(tree)$ ;
4:    $f^i, reward = \text{Simulate}(\mathbf{a}^i)$ ;
5:   if  $f^i > \lambda$  then
6:      $\mathcal{A}_{desired} \leftarrow \mathbf{a}^i$ ;
7:   end if
8:    $\text{Backpropagate}(tree, reward)$ ;
9: until Terminate
  
```

A. L-system Semantic Interaction Sequence and Tree

L-system is a string rewriting mechanism and is perfectly suitable to represent the cascades of **interactions** due to its strength in describing sequential branching behaviors [5], [14]. Additionally, L-system expands the string in a hierarchical fashion which fits the traditional MCTS usage. In general, L-system is expressed as $\mathcal{G} = \langle \mathcal{V}, \Omega, \mathcal{P} \rangle$ with \mathcal{V} being a set of alphabets for symbolization, Ω being the axiom, i.e., a word, and \mathcal{P} signifying a set of rewriting rules. The produced L-string is composed of a sequence of alphabets where the alphabets reveal e^i and their order reveals t^i , forming the alternative term \mathbf{l}^i that embeds both the **elements** and **interactions**. Even though the L-system, \mathcal{G} , completely depends on the application and cannot be generalized into a common form, we involve one commonly used rewriting rule to split the interaction sequence into multiple branches with brackets. That is, the sub-string wrapped in the brackets belongs to a single **interaction** path and would not affect other branches, resulting in the form of $str_0[str_1][str_2] \cdots [str_i]$ where str_1 to str_i are the sub-strings branching from the last alphabet of str_0 .

To derive an L-system from the given application, there exists no single unique answer and the derivation might greatly vary under different objective goals. This work intends

to show its deduction ability for coherent and hierarchical exploration with three exemplar L-systems based on their respective domain knowledge in Section V, but we do not intend to develop a general L-system derivation for cascade events. Although the L-string itself can well represent the **elements** and **interactions**, we would like to further deduce it into the tree form, named as the interaction tree, to describe the spatio-temporal content in a hierarchical fashion. Such a tree not only provides a more intuitive visualization of the **interactions** but also eases the later process when performing inverse embedding (please refer to Section IV-B for more details). To turn an L-string into a tree, the process starts from the first alphabet and iteratively appends the next alphabet as a child to the current alphabet. The tree splits into branches when the bracket is met, with the sub-string in the bracket forming one single sub-tree. The deducing example can also be found in Section V for different applications.

B. Inverse Embedding

There are two common ways to derive the corresponding actions from a distribution of $p_w(\mathbf{x}^i)$: one is the forward derivation, and the other is the inverse postulation. Taking the one-shot-all-sink billiard set-up as an example, forward derivation first specifies all possible configurations, including positions and the cue strike conditions, and then runs physical simulation to determine the successfulness, i.e., estimate the probability; on the other hand, inverse postulation uses the collision physics to estimate the possible travelling path of each ball to fulfill the pocketing requirement while deriving the corresponding positions and initial conditions at the same time. The forward, however, necessitates the exploitation of almost infinite space, resulting in computational intensity. On the other hand, the inverse could also be impractical since all path trajectories would be intractable if complex physics involved. As a practical solution, coherent sampling algorithms, such as Markov chain Monte Carlo (MCMC) and Monte Carlo Tree Search (MCTS), are proposed to effectively reduce the amount of exploration. However, a pure forward solution seems to be inefficient when the domain knowledge can be applied to greatly reduce the explored space, resulting in a more plausible and efficient algorithm. For example, the possible position of an object ball can be restricted to a sector area in front of the pocket, while the forward strategy requires to sample all possible positions on the entire table. Therefore, we propose to incorporate the inverse concept into MCTS, which is in general a forward solution, as a hybrid method to model $p_w(\mathbf{x}^i)$. More specifically, inverse embedding, derived based on the expertise, is performed to reduce the sample space of \mathbf{x}^i for any given L-string. However, the amount of domain knowledge and the level of procedural factors taken to postulate the forward parameters is a trade-off between computational complexity and sampling efficiency. This work intends to show its strength in searching space reduction and sampling efficiency enhancement with three domain-specific inverse postulations shown in Section V, but it is not our intention to provide a general inverse solution for constraint animations.

C. L-system Monte Carlo Tree Search (L-MCTS)

Traditional MCTS expands the search tree with four stages of selection, expansion, simulation, and backpropagation as defined in the following:

- **Selection** selects successive children from the root to a leaf that is not yet fully explored.
- **Expansion** appends children with not yet investigated states to a leaf.
- **Simulation** performs a complete playout from the current until a result or predefined state is achieved.
- **Backpropagation** updates the probability distribution based on the evaluated objectiveness and is always performed right after the simulation operation.

However, some modifications are incorporated into the proposed L-MCTS in order to better fit our applications. First, we only perform the simulation when all parameters are decided, including l^i , x^i , and z^i , while traditional MCTS performs simulation at any node with the following moves decided randomly. More specifically, we keep performing selection and expansion until simulation is available. The key to such modification is that the number of parameters to sample is relatively constant in our case, compared to traditional cases in which the number of moves to finish the game is usually difficult to predict. Second is the node-choosing strategy. Traditionally, MCTS performs selection based on a designed metric to balance exploration and exploitation. The most popular metric is the Upper Confidence Bounds to Trees (UCT) [30], which in general can be formulated as $c^j = \frac{e^j}{n^j} + \sqrt{2} \sqrt{\frac{\ln(n^{j\uparrow 1})}{n^j}}$, where j stands for any node from the tree, n^j stands for the simulation count of the j -th node, $n^{j\uparrow 1}$ stands for the simulation count for the parent of the j -th node, and $\frac{e^j}{n^j}$ implies the empirical mean. An intuition is that for an untraversed node, the UCT evaluates to infinity, and therefore, traditional MCTS chooses to perform expansion prior to selection if there are unexplored states. However, as we are sampling parameters from continuous spaces, such behavior would never lead the search to converge. Hence, we slightly modify the definition of **Selection** and **Expansion** into "stepping down the hierarchy by picking one child", while the former only picks the existing child and the latter appends a new one. Also, a new strategy is proposed for deciding whether selection or expansion is performed based on $C^j = \sum c^{j\downarrow 1}$, where $c^{j\downarrow 1}$ represents all children of the given j -th node. The concept behind the strategy is a trade-off between exploitation and exploration, i.e., we choose selection if the current parameter sets are promising enough to generate desired results; otherwise choose expansion for exploring other possibilities. Technically, we choose to perform selection if $C^j \geq C_\alpha$. Such a threshold C_α might vary under different search scenarios, and hence, it is decided empirically for different applications proposed in this paper. We introduce the experiment involved to decide the threshold in Section VI-A1

Another difference is that the simulation result of L-MCTS ranges between $[0, 1]$ to imply the preference instead of simple win or lose. More specifically, we have $\Delta e = \text{Max}(1 - \frac{|\lambda - f(\mathbf{a}^i)|}{\beta}, 0)$ and add Δe to e^j for nodes on the path from root to leaf when performing backpropagation. λ reflects the objectiveness, but its value is application-dependent and not

normalized to the range $[0, 1]$. Therefore, this work uses β to normalize the application-dependent objectiveness deviation, and its value is decided empirically.

The process to sample \mathbf{a}^i can be separated into three phases: L-string expansion, spatial positioning, and activating initialization while each distributes in the search tree from top to down respectively, and different actions are taken to sample a new parameter set during different phases.

- **L-string expansion** aims to sample the L-strings by inserting alphabets into the string iteratively. For any given node in this phase, the number of its child is one plus the number of its L-string variants with one rewriting rule applied, where the additional node represents to step down to spatial positioning with the current string. Since the node number in one layer is finite in this phase, we follow the traditional strategy to select a child with the largest c^j . Also, to avoid the string growing infinitely, we simply restrict the string length to n_{max} , which depends on the search goal and application.
- **Spatial positioning** aims to place the elements that come from L-string expansion into the space. Any given node in this phase represents the position of one corresponding element, where its expansion is conditioned by the inverse embedding mechanism. The propagation direction can be either forward, i.e., from root to leaves given the interaction tree, or inverse vice versa. In general, both approaches are acceptable, but the better one is application-dependent.
- **Activating initialization** aims to decide initial conditions to trigger the spatio-temporal content. Practically, this part is relatively simple as the initial conditions are independent with each other in common cases. Thus, it results in a single-layer that samples all parameters in once. Furthermore, the phase could be optional if a sequence does not require any explicit triggering action.

D. Objective Function

The objective function directly defines the search goal of L-MCTS, and it could be either objective or subjective and usually fully application-dependent. Even though our L-MCTS does not constrain it to any specific form, it should be positively correlated with the number of interactions, i.e., the length of the sequence, to facilitate coherent exploration of MCTS. As examples, we show the possible derivation of objectiveness by incorporating objective metrics, such as the animation elapsing time, i.e., application-independent, as well as subjective metrics, such as the visual complexity, i.e., application-dependent in Section V. Furthermore, we also demonstrate a procedure to derive the subjective evaluation based on questionnaire postulation of user preferences in Section V-B4 and Section VI-A2.

V. APPLICATIONS

Our L-MCTS is designed to be a general framework of massive production for various applications. To demonstrate its generality, we apply it on three different applications, including one-shot-all-eliminate breakout clone, chain reaction, and one-shot-all-sink billiard. The following sections introduce

them respectively along with their specific L-system, inverse embedding, activating initiation, and objective function derived from their unique domain knowledge. For brevity, we omit the unit of length if it is measured in meter. Also, the details of each application are provided for reproducibility, while some design choices are omitted due to length issue. Readers can refer to the supplemental Website¹ for more details.

A. One-shot-all-eliminate Breakout Clone

The breakout clone is a famous video game that a player controls the moving paddle located at the bottom to make the ball hit the bricks and destroy them. The ball bounces when hitting either the paddle, a brick, or the top, left, and right wall, and the player loses when the ball hit the bottom. As shown in the left of Fig. 4, we derive a simple toy example from it as our first validation attempt. We let the ball destroy all bricks without another bounce from the paddle, i.e., once the ball is emitted from the paddle, it never lands but destroys all bricks aligned horizontally.

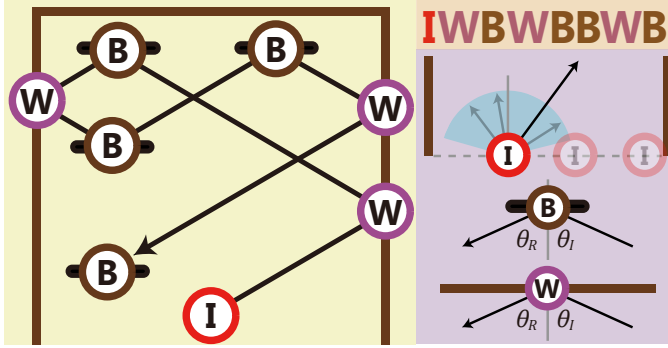


Fig. 4. This shows an example of our one-shot-all-eliminate breakout clone with its representative L-string. In addition, the inverse embedding technique according to the reflection rule is also visualized in the right bottom.

1) *L-string Expansion*: Based on the bouncing behaviors, we can semantically categorize them into two bouncing types, wall bounces, W , and brick bounces, B . We also use I to denote the emission position of the ball. The system always begins with the axiom, IB , and extends the level by iteratively inserting consecutive bounces into the current string:

- The **wall bounce** inserts a W into any position of the current string, except for the very front and the very back.
- The **brick bounce** inserts a B into any position of the current string, except for the very front.

An example string as well as its possible level is shown in Fig. 4.

2) *Spatial Positioning*: Accordingly, the face normal at the bouncing point is the only affecting factor of ball movement, i.e., the outgoing direction is based on the reflection of perfect elastic collision where the incidence angle equals the reflection angle. To decide the element positions, we start from the root of the interaction tree and propagate to its descendant in a forward order. More specifically, we decide both the element position, p^i , and outgoing direction, \vec{v}^i , for the i -th node based on the element type:

- **I**: p^i represents the emission point randomly chosen along the paddle moving line, while \vec{v}^i is the emission direction whose angle between the vertical axis is sampled within the range of $[-65^\circ, 65^\circ]$.
- **W**: p^i represents the bouncing point of the wall located as the first intersection of the ray $p^{(i-1)} + \vec{v}^{(i-1)}$ to the three walls, while \vec{v}^i is the bouncing direction derived from the reflection rule, i.e., to have the reflection angle θ_R equal to the incidence angle θ_I .
- **B**: p^i represents the bouncing point of the brick randomly chosen along the outgoing path of the previous node, while \vec{v}^i is the bouncing direction derived from the reflection rule.

After processing all nodes, we place the paddle as well as other bricks with their center locating at p^i . We demonstrate the spatial positioning mechanism in the right bottom of Fig. 4.

3) *Initialization*: To initiate the break-out clone, the ball is emitted into the scene from the center of the paddle with an outgoing direction. Since the paddle position and outgoing direction is already determined in the spatial positioning phase, there are no extra factors sampled in the initialization phase.

4) *Objective Functions*: We evaluate the breakout clone level in terms of difficulty to figure out the correct bouncing path to eliminate all bricks in once for the given level. It is harder to plan the bouncing path if the two bouncing points are farther, i.e., the difficulty should be proportional to the distance between two bouncing points. If the number of bounces increases, the difficulty increases accordingly, i.e., the difficulty should be the accumulation of difficulty at each bouncing. As a result, we can express the difficulty as $f_d = \sum \alpha^i \cdot (1 + 0.2 \cdot \frac{l^i}{l_{diag}})$ for $i \neq 0$, where l represents the diagonal length of the scene, and α is the bouncing weight which is set to 15 for W and 10 for B due to the fact that a brick is smaller and harder to predict its bounce. Furthermore, we have $\beta = 0.25 \cdot \lambda$ to normalize the reward.

B. Chain Reaction

Chain reaction is the most classic application of spatio-temporal content because of its cascading events to propagate interactions from one component to another. In general, components involved in chain reaction can be anything as long as it makes cascading propagation continue, and the allowable interactions between components are totally implementation dependent. Here, we choose several commonly seen components, such as balls and dominoes, to produce chain reaction animations of targeting visual complexity. As shown in Fig. 5, we restrict the interaction path to a 2D plane in general, while the path might split into multiple ones where each lies on a 2D plane paralleling to each other.

1) *L-string Expansion*: We denote the minimum unit to propagate the events as a component and semantically derive the L-system representations by assigning an alphabet to each component while the production rules are then carefully designed for plausible propagation. To demonstrate our framework, we select 5 different types of components visualized in Fig. 6. The first component, B , is composed of a ball placed on a slope, and the ball rolls down the slope. Second, we

¹Website: <http://graphics.csie.ntust.edu.tw/pub/LMCTS/>




Fig. 5. This demonstrates an example of our chain reaction set-up with its representative L-string, where the interaction paths are designed to lie on 3 2D planes marked with red, green, and blue, respectively. The action starts from the ball that rolls down the slope, B , striking 3 balls on the platform, L , into the cups, and activating their respective pulley sets, P . While two of them drop into the cup to end their sequences, the other topples the domino run, D , which ends up with the last domino falling into the cup, C .

have the domino set, D , that is straightly placed along the x -axis and can be toppled from the first domino. Third, we have the pulley set, P , with one side being the cup and the other being the baffle, while the baffle is heavier than the cup to block the ball from rolling down the slope. In other words, the pulley is only triggered when something falls into the cup and makes the ball start rolling down the slope because the baffle gets removed. Additionally, a platform with multiple balls placed on it, referred as L , is included to perform the branching behavior. More specifically, balls on the platform can be struck into different directions and split the interaction path into multiple ones where each lies on a new 2d plane. Lastly, we have the single cup, C , to be the end of the chain reaction, i.e., the cascading events end when the previous component falls into the cup. For each component, we propose a corresponding production rule to insert a new instance into the current set-up as following:

- The **dominoes**, D , can only be toppled by the rolling ball or another domino set and cannot trigger multiple balls on the platform, which means D can only be inserted after B or D and not before L .
- The **pulley set**, P , and the **ball on slope**, B , are grouped together as PB . Since the pulley can be triggered by falling components including the domino and ball, PB must be inserted after B , D or L .
- The **platform with multiple balls**, L , can only be triggered by the rolling ball and only be used to trigger the pulley set, i.e., we insert L after B and before P . Also, due to the branching behavior of L , the interaction path splits into n paths with n balls placed on the platform, with one sub-path being the original sub-string after the insertion point and others being initialized to be C . To simplify the set-up, we restrict n to be 3 and the branching behavior can only be inserted to the main branch as illustrated in the top left of Fig. 6.
- The **cup**, C , can only be placed at the end of the string and no production rule inserts C into the string directly.

The chain reaction always begins with BC and iteratively applies the insertion rules to extend the interaction paths.




Fig. 6. This visualizes the 5 components involved in our chain reaction application, where the V_L , V_B , and V_P decide the placing directions along the x -axis. Also, we define the **reference input**, r_{in} , and the **reference output**, r_{out} , for each component to align two components with additional offset δ_x and δ_y for the inverse embedding mechanism.

2) *Spatial Positioning*: To increase the behavior variety, we involve additional control parameters for each component as below.

- The **dominoes** are parameterized by the domino number, n_D , within the range $[4, 10]$, and the interval between each domino, l_D , within the range $[0.2, 0.3]$.
- The **pulley** is parameterized by the distance between two pulleys, l_P , within the range $[2, 5]$, the distance between the pulley and its carrying object, l_c and l_b , within the range $[0.5, 1]$, and the facing direction v_P to be either $\{-1, 1\}$ to represent the alignment along the x -axis.
- The **ball on slope** is parameterized by the slope length, l_B , within the range of $[1, 5]$, the elevation angle between the slope and the x -axis, θ_B , within the range $[26^\circ, 30^\circ]$, and the facing direction v_B to be either $\{-1, 1\}$.
- The **platform with multiple balls** is parameterized by the facing direction v_L to be either $\{-1, 1\}$. The first ball would be put on the position which offsets $4 \cdot r \cdot v_L$ from the center of the platform along the x -axis, where r represents the ball radius, and the followings are put on the position which offsets $\frac{4}{3} \cdot r \cdot \alpha_L$ from the center along the x -axis, where α_L is -1 for the second ball and 1 for the third ball.
- The **cup** is not parameterized.

Generally speaking, we arbitrarily choose some parameters to have an adequate sampling range for compact illustration. For example, n_D is chosen to have a reasonable length of domino toppling; l_P is selected to make the pulleys not too far away. At the same time, the others are decided empirically to provide acceptable success rate. For example, l_B and θ_B are chosen to provide adequate energy applied to the ball for triggering the next component. Due to the length limitation, those parameters, such as mass and dimensions, chosen for producing the results are summarized in Appendix A.

To set up a chain reaction that can successfully propagate the interactions to the end, it is difficult not to perform trial-

and-error even for the experts because the involving physics could be extremely complex. In fact, the experts usually place components based on the intuition in the first place and make small modifications according to the simulation result. Following such a concept, we define r_{in} and r_{out} for each component to represent the point of **reference input** and **reference output**, visualized in Fig. 6, where the former one indicates the point to successfully trigger the component, and the latter one indicates the point that the component propagates the interaction to its descending path at. When aligning two reference points, in order to avoid component overlapping as well as to handle the unpredictable behavior of physics, we add an extra 2D random offset, (δ_x, δ_y) , both sampled in the range of $[0.2, 1]$. The overall workflow of spatial positioning begins from the root and processes nodes in a forward order, while each component is first parameterized and placed to the scene by aligning its r_{in} to previous r_{out} using the above process. The only exception is the root position which is randomly sampled from the 3D space.

3) *Initialization*: The initial conditions determine how much energy the chain reaction system has in the beginning. While we always have the first component to be a B according to our production rules, the initial conditions fully depend on the initial height of the first ball. More specifically, we sampled a distance ranging between $[10, 20]$ that represents the height difference between the ball and the slope.

4) *Objective Function*: Chain reaction is evaluated based on two different metrics: f_t simply reveals the elapsed time from begin to end, and f_c evaluates the visual complexity of the set-up. f_t helps to validate the effectiveness of the proposed framework, and f_c generally involves subjective opinions and consensus. To the best of our knowledge, no existing literature has ever formulated and postulated it. Therefore, this work formulates it empirically based on the questionnaire methodology introduced in Section VI and proposes to model f_c as

$$f_c = \sum H(i) + \alpha_g G(\cdot), \quad (3)$$

where $H(i)$ evaluates the component-wise complexity for the i -th component, $G(\cdot)$ evaluates the layout complexity as a global term, and $\alpha_g = 0.4$ is the empirical weight to balance between the local and global term. The component term, $H(i)$, is defined as $H(i) = c^i \cdot k^i$, where c is the scaling factor to normalize the complexity contribution into the same range for each type, and k is the complexity contribution. We have $k = \frac{K}{R_{max}}$ where K is the value of the control parameter which has the most dominant affection in complexity, and R_{max} is the maximal allowable value of the control parameter. Accordingly, we set c and K as follows:

- **B** has K to be the slope length, and sets $c^i = 1.00$.
- **D** has K to be the domino number, and sets $c^i = 2.75$.
- **P** chooses K to be the distance between two pulleys, and sets $c^i = 1.22$.
- **L** chooses K to be the branch number, and sets $c^i = 2.81$.
- **C** has $K = 1$ because the cup is not parameterized, and sets $c^i = 0.1$.

The global term is modelled as $G(\cdot) = w_b \cdot g_b + w_d \cdot g_d + w_v \cdot g_v$, where g_b represents the tree balance, g_d represents the depth of

the interaction tree, g_v represents the number that movement direction changes, and w_b, w_d , and w_v are the weights of each term set to be 1.88, 2.25, and 1.88, respectively. The intuition behind the balance term is to reveal the number of interactions happened simultaneously because the system looks more complex when there are more interactions happened in the same time. Specifically we have $g_b = \sum \frac{d_{avg}^j}{d_{max}^j}$, where d_{avg}^j and d_{max}^j represents the average and maximal branch depth at the j -th branching point where the branch depth is defined as the depth from the branching point to either the leaf node or the next branching point. All the above mentioned weights are decided based on the user study discussed in Section VI-A2. The only exception is cup's weight set to 0.1 because it remains static and unchanged throughout the animation. Also, we set $\beta = 2 \cdot \lambda$ to normalize the reward.

C. One-shot-all-sink Billiard Layout

Billiard is a well-known sport that strikes balls into the pockets with a cue. Among such a category, there is a special variant that challenges the player to sink all object balls in only one shot, which is referred as the one-shot-all-sink billiard. To design a plausible one-shot-all-sink layout such as the one shown in the right top of Fig. 7, we define several terminologies at first hand to systematically explore the domain. Ghost ball is an abstract concept that we use to represent the event of ball-to-ball collision, and a kiss shot or a borrow shot is referred to the pocketing event while the object ball from the former shot directly goes into the pocket and the one from latter shot collides with other balls before pocketing. We apply the proposed content generation framework on Pool Ace, which is the commercial game product provided by the International Gaming System (IGS), with Unity 2019.

1) *L-string Expansion*: Conceptually, experts use the collisions among the cue, objects, cushions, and pockets to describe billiard interactions. Therefore, to semantically derive the L-system with the billiard application, we conclude two collision types, which are the ball-to-ball collision denoted as G and the ball-to-cushion collision denoted as C , as well as two pocketing strategies, which are the kiss shot referred as P and the borrow shot referred as X . Also, we involve some alphabets to represent the ball position, which are I, S , and O for the begin position of the cue ball, the end position of the cue ball, and the begin position of the object ball respectively. We design the production rules that insert shot combinations, including the cushion and kiss shots, into the existing trajectory, while the borrow shot is always transformed from the kiss shot in specific situation. More specifically, the production rules are listed as below:

- **Cushion** inserts an extra cushion shot, C , on any moving trajectory, i.e., after I, O, C , or G but not between GO .
- **Kiss** inserts an extra kiss shot, GOP , on any moving trajectory and results in a splitting branch. If the insertion happens before any kiss shot, i.e., P , change it into X as the shot becomes a borrow shot.

The system always begins with the axiom, $IG[S][OP]$, and extends the set-up by iteratively applying one of the two production rules each time. When transforming the L-string

into the interaction tree, O is omitted because GO always show up together. An example string and its possible layout is shown in Fig. 7, with its deducing interaction tree placed on the right.

2) *Spatial Positioning*: In Pool Ace, all the balls are forced to remain on the table and no bouncing happens during the play, so we design the inverse embedding techniques based on such assumption to postulate the initial positions of the cue and objects. We first decide the end position and propagate to the rest of the balls in a backward order, due to the fact that pockets are constrained to fixed positions on the table. That is, the positions are decided by processing each node from leaves to root, given the interaction tree transformed from an L-string. To decide the position of leaf nodes, we randomly sample from the 6 pockets for P and X , and from the circle region centered at its sibling for S , with the radius set to $\frac{1}{3}$ of the table diagonal. On the other hand, positions of the non-leaf node are decided based on the inverse embedding mechanism to narrow the sampling space into a more reasonable subspace. While precisely predicting the ball trajectory is computationally expensive, we perform inverse embedding by simply assuming particle balls and elastic collisions without frictions to avoid costly computation. More specifically, position of a non-leaf node is sampled based on the following two constraints. The first one is the circle constraint that is only applied when there exists two child nodes. One circle is constructed by connecting its two children as the diameter, and the node is constrained to the circle. The second one is the direction constraint that is applied to each individual child node separately. A fan area is constructed by spanning θ_p degrees to two sides of the direction \vec{n} , and the node is constrained to be inside the area, where θ_p and \vec{n} vary for the different child node:

- **P/X**: \vec{n} represents the angle bisector and $\theta_p = 22.5$ for pockets on the corner and $\theta_p = 45$ for others.
- **G**: \vec{n} represents the outgoing direction passing the child node from the circle that G is sampled from and $\theta_p = 5$.
- **C**: \vec{n} represents the reflection path and $\theta_p = 0$.

When multiple constraints are applied, the node is sampled from the intersection of all constraints. The mechanism is demonstrated in Fig. 7 with the given L-string and the final spatial distribution is shown on the right top.

3) *Initialization*: The initial condition refers to the cue hitting behavior and is defined as $H(\vec{f}, p)$ in our billiard application, where \vec{f} represents the directional applied force and the p reveals the hitting point on the ball surface. In general, the cue ball draws a curved trajectory before the first collision if the applied force does not pass the ball center, while the rotation velocity usually vanishes quickly due to the friction, leaving the movement to be straight on the remaining path.

4) *Objective Function*: After examining billiard instructions and books [31], [32], experts generally evaluate the difficulty by considering the path planning as well as the striking execution. Furthermore, the striking effect is exponentially waded off by the friction, so the initial cue hitting effects become negligible for later collisions. Therefore, the difficulty metric, f_d , is empirically designed to be a summation of the




Fig. 7. This shows an exemplar of our one-shot-all-sink billiard application as well as its spatial positioning process: (a) the position of leaf nodes are sampled from their candidate space, (b) the ghost ball G with two children is placed on the intersection of the circle and the fan areas. (c) the cushion C with one single child is constrained by the fan area and the table edge. (d) the ghost ball G with one cushion child is placed on the intersection of the circle and the reflection direction that satisfies the equation of $\phi = \theta$.

striking term and the collision term:

$$f_d = \alpha G(\cdot) + \sum H(i), \quad (4)$$

where $G(\cdot)$ denotes the cue-hitting difficulty, $H(i)$ denotes the collision-wise difficulty measurement for each collision, and α is a user specified weight to balance the two, set to 0.5 in all billiard experiments in this paper.

For the striking term, we assume that spin and power are the two most effective factors to predict the ball trajectory and have $G(\cdot) = \|\vec{f}\| l_G^2$, where $\|\vec{f}\|$ is the applied power in newton (N), and l_G is the distance between the center and hit point in cm . For the collision-wise term, the intuition is that the more consecutive collisions happen before a specific collision or the longer the distance a ball travels from the previous collision, the harder it is to predict and execute the desired process. Moreover, while the exiting direction deviates more from the entering, it is also harder to aim. All three factors pull together to determine the measurement of the collision i as $H(i) = l_H^i n^i A(\cdot)$, where l_H^i represents the traveling distance from collision $i-1$ to i , n^i represents the accumulated number of consecutive collisions before i , and $A(\cdot)$ is measured based on the angle information depending on the following collision types:

- **Cushioning** happens between a ball and a cushion. $A(\cdot)$ is measured as $\sin(\theta_I^i) \text{Max}(\frac{\tan(\theta_I^i)}{\tan(\theta_R^i)}, \frac{\tan(\theta_R^i)}{\tan(\theta_I^i)})$, where θ_I^i and θ_R^i represent the incident and reflective angles. In general, it is more difficult and unpredictable when θ_I^i is larger and the difference between θ_I^i and θ_R^i is larger.
- **Colliding** happens between two balls. We use the cut angle θ_c^i , which is the angle between the incoming direction and the reflection axis, and the tail angle θ_t^i , which is the angle between the outgoing direction and the extension of the incoming direction, as $\frac{1}{\cos(\theta_c^i) \cdot \cos(\theta_t^i)}$ because it is easier to aim while the enter and exit directions are near parallel.
- **Pocketing** refers to a ball entering a pocket. We use the pocket angle, θ_p^i , which represents the angle between the angle bisector of the pocket and the entering direction, to measure $A(\cdot)$ as $\frac{1}{\cos(\theta_p^i)}$, due to the intuition that it is

easier to hit the corner of the pocket and pop out if θ_P^i is larger.

The lengths and angles involved in the collision-wise difficulty measurement are visualized in Fig. 8. We set $\beta = 0.25 \cdot \lambda$ to normalize the reward.




Fig. 8. This illustrates the collision-wise difficulty metric calculated based on the traveling lengths and collision angles where the lengths l_H^i is the distance between parent and node collisions, the incident angle θ_L^i and the reflective angle θ_R^i are used for cushioning, the cut angle θ_c^i and tail angle θ_t^i are selected for ball-to-ball collision, and the pocket angle θ_P^i is chosen when pocketing.

VI. EXPERIMENTS AND RESULTS

After implementing three targeting applications, we conduct several experiments and user studies to decide their applied parameters. Furthermore, we compare our proposed method to three baselines in the aspect of content generation with the targeted goal. Finally, since objective functions are important and abstract, we conduct user studies to validate their effectiveness. Due to length limitation, readers can refer to the supplemental Website¹ for more details.

A. Parameter Selection

Generally, our applications have several user-tuning parameters of either an exploration/exploitation trade-off or user's subjective preferences, which cannot be directly decided. As a result, we conduct experiments and user studies for proper selection.

1) *Selection Threshold*: Our L-MCTS chooses to perform **selection** or **expansion** depends on whether C^j is greater than the selection threshold C_α . The choice of C_α itself is, however, a trade-off between exploration and exploitation. A smaller C_α makes the system focus more on the previously explored results while a larger one tends to exploit other new parameter spaces. To decide the threshold, we first evaluate the search with different C_α 's to figure out its actual impact on the search task. More specifically, we perform the search tasks with different C_α 's and evaluate the desired rate ρ_γ , which is defined as the number of desired instances out of the total number of searching. A desired instance is defined as \mathbf{a}^i that satisfies $\frac{f(\mathbf{a}^i) - \lambda}{\lambda} < \gamma$, where the left term represents the error percentage and the right is the error tolerance. We test with 5 different C_α 's, which are $\{2.0, 2.5, 3.0, 3.5, 4.0\}$, on the chain reaction with its error tolerance γ set to 0.05. For each, we search 5000 iterations, and set the maximum string length n_{max} to be 8 and the desired goal λ to be 8 seconds. The left of Fig. 9 shows the desired rate under different C_α 's across the search, while $C_\alpha = 2$ generally fails to perform the search, and $C_\alpha = 3$ outperforms others when the iteration number increases. The right of Fig. 9 shows the exploration rate across

the search, which is defined as the none-repeated rate of the parameters we sampled from the continuous space. In general, such rate indicates the number of explored parameters and decreases more when the C_α is smaller. Accordingly, we choose $C_\alpha = 3.0$ for the chain reaction due to its significant performance gap against others and its acceptable exploration rate after 5000 iterations. Similarly, we apply the same process to set C_α to 3.0 and 1.85 for the breakout and billiard, respectively.




Fig. 9. We visualize the desired rate (left) and the exploration rate (right) under different C_α while searching. The iteration number is denoted along the x -axis. For the y -axis, we have the desired rate, defined as the number of instances whose error percentage less than 0.05 out of the total number of searching, in the left, and have the exploration rate, defined as the none-repeated rate of the parameters we sampled from the continuous space, in the right. The experiment result of different C_α is visualized in different colors denoted in the legend.

2) *Subjective Weightings*: In Section V-B4, we select several factors to be our indicator of the visual complexity and formulate them into Eq. 3 as our complexity metric. However, the weight of each factor remains undetermined because they are subjective preferences. Therefore, we designed a user study to ask users to sort them based on visual complexity from the least dominant to the most dominant. It was conducted on 16 subjects, aged from 20 to 29. Among them, there were 15 males and 1 female, and also there were 7 majoring in computer graphics and 2 having experience in animation design. At the beginning of the user study, we showed 6 chain reaction set-ups to the users and explains the meanings of the 7 factors. Factors of the 4 components, including B , D , P , and L , are explained as the definition listed in Section V-B4. Besides, g_b is given as "whether the branches split from same origin are equally long or not", g_d as "the longest path in the scene", and g_v as the "direction change among the set-up". Then, the users were asked to perform the sorting via the factor dominance on visual complexity according to their own subjective opinions. The given orders are treated as the **raw scores**. Since the scope of g_b , g_d , and g_v involved the the entire content generation, we categorized them as **global factors**, for others being the local factors, and compute the orders among the local/global factors to acquire the **local scores** and the **global scores**. We demonstrate the scores under different domains in Fig. 10 along with their averages and standard deviations. We use the average of **local scores** and **global scores** as the weights, while the corresponding weight of B and P is further multiplied with 0.5 to half the providing contribution since they are always bounded together, and the weight of C , which is not involved in the user study, is directly set to 0.1 since it remains static during the chain reaction.

Later, Section VI-C validates their effectiveness.

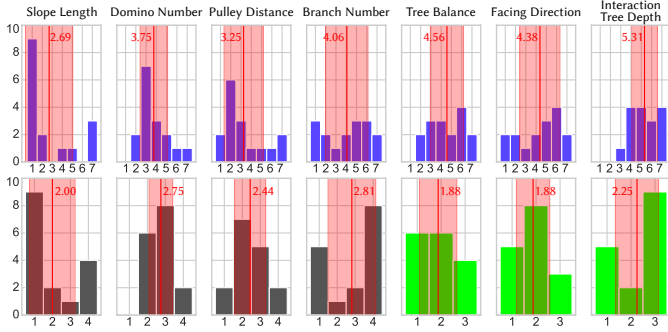


Fig. 10. This shows the histograms of our subjective weighting study for the chain-reaction visual complexity metric along with their average (the red line) and standard deviation (the red zone). Blue histograms represent the *raw scores*, while the black ones represent the *local scores* and the green ones represent the *global scores*.

B. Efficiency Enhancement in Content Generation

Since the L-MCTS is claimed as the forward solution incorporating the inverse mechanism, we evaluate the system in two aspects:

- 1) **The choice of the sampling strategy:** We compare the MCTS with the random sampling to show the strength in systematical and coherent exploration.
- 2) **The effectiveness of the inverse embedding:** We compare the search efficiency of the framework with inverse embedding against the one without the mechanism, i.e., randomly placing elements into the scene.

Accordingly, we construct three baselines based on the combinations of (1) and (2) and denote them as (a) **random sampling with inverse embedding (RSIE)**, (b) **MCTS without inverse embedding (MCTS)** and (c) **random sampling without inverse embedding (RS)**. Technically, random sampling also samples a solution based on the 3 phases of L-string expansion, spatial positioning, and initialization, while L-string expansion samples one single string from all possible L-strings restricted by the maximum string length of n_{max} , and the rest of two phases performs random sampling based on the selected L-string. Those, that do not involve inverse embedding, choose elements' position from the entire space, while in chain reaction, the infinitely large space is restricted to a cube with a dimension of $[100, 100, 2]$. We apply all 4 methods to our three exemplar applications, with the targeted goal and maximum string length configured as following:

- **One-shot-all-eliminate Breakout Clone:** we set target difficulty to 75, $C_\alpha = 3.0$ and $n_{max} = 6$.
- **Chain Reaction:** we set desired time duration to 10 seconds, $C_\alpha = 3.0$ and $n_{max} = 10$.
- **One-shot-all-sink Billiard:** we set target difficulty to 1000, $C_\alpha = 1.85$ and $n_{max} = 6$

We use the same C_α for both our L-MCTS and the MCTS. For each method, we randomly generate 5000 sets and record the number achieved under the targeting error percentage, γ , of 0.05, 0.025, and 0.01. As shown in Table I, our L-MCTS outperforms others with a significant gap. Accordingly, we can conclude with two folds. On one hand, 3 compared methods show low searching effectiveness to demonstrate the difficulty

TABLE I
THIS SHOWS THE STATISTICS OF GENERATION EFFICIENCY UNDER SPECIFIC ERROR PERCENTAGES, γ , FOR 3 DIFFERENT APPLICATIONS WHILE COMPARING OUR L-MCTS AGAINST THREE OTHER BASELINES.

	Breakout Clone			Chain Reaction			Billiard		
	$\gamma = 0.05$	$\gamma = 0.025$	$\gamma = 0.01$	$\gamma = 0.05$	$\gamma = 0.025$	$\gamma = 0.01$	$\gamma = 0.05$	$\gamma = 0.025$	$\gamma = 0.01$
L-MCTS	4519	4347	683	2515	1468	501	2963	2775	854
RSIE	195	101	24	27	10	3	13	6	3
MCTS	4	2	2	0	0	0	4	2	1
RS	9	3	2	0	0	0	1	1	0

to sample goal-fulfilling instances while L-MCTS can enhance this search efficiency using coherently hierarchical exploration. On the other hand, the inverse embedding technique shows significant impact on improving the search efficiency, but the choice of the forward sampling strategy also plays an important role to produce massive amounts of desired results.

C. Chain Reaction Complexity Metric

To evaluate whether the proposed subjective metric is capable of revealing the viewers' opinion, we conducted a user study that asked the subjects to grade the chain reaction scene based on the visual complexity. First of all, we categorized the scenes into 5 complexity levels whose mean complexity is $32.6 \cdot 1.75^i, i \in \{1, 2, 3, 4, 5\}$. The category was designed to be equally distributed within the minimum and maximum complexity, which were both chosen based on authors' experiences. We generated 2 scenes from each level with the error percentage threshold γ set to 0.05 as the *reference scenes*. Also, we generated 2 scenes for each level, resulting in 10 scenes in total, and denoted them as the *evaluation scenes*. Due to the length limitation, here, we only show one scene for $i = 1, 3, 5$ from the *reference scenes* and the *evaluation scenes* in Fig. 11. During the user study, we first showed the *reference scenes* to the subjects as well as their designated visual complexity levels. Then, all participants were asked to assign their subjective complexity level for all *evaluation scenes* arranged in a random order in a 5-point Likert scale manner, where 1 represented the least complex and 5 represented the most complex. We conducted the user study on 16 subjects who were aged between 20 and 29, with 15 males and 1 female, and there were 7 majoring in computer graphics and 2 had experience in animation design. We sum up all user assigned levels for each scene as its complexity scores, shown in Fig. 12. It is obvious to see the user assigned level is highly correlated to the designated level. To further justify the assumption, we calculate the Spearman's rank correlation coefficient [33] between i and \bar{i} , for \bar{i} being the user assigned level, and acquire $\rho = 0.9833, p \approx 0.0$. The detailed equation is provided in Appendix B. The interpretation of $\rho > 0$ shows that the user assigned level is positively correlated to the designated level. In addition, a $|\rho|$ closer to 1 indicates a higher degree of correlation, and $|\rho| > 0.7$ is usually seen as a strong correlation evidence [34]. Also, the p-value indicates the probability to observe such result under the null hypothesis being $\rho = 0$, and since $p \approx 0.0 < 10^{-3}$, it is statistically significant to conclude that the proposed complexity metric can reveal the viewers' opinion.




Fig. 11. This shows several examples of the *reference scenes* and the *evaluation scenes* for the complexity study of the chain reaction.




Fig. 12. This shows the user assigned complexity levels of chain reaction scenes, while the designated complexity level is highlighted in red in the legend.

D. One-Shot-All-Sink Billiard Layout Difficulty Metric

We have also conducted a user study to validate the difficulty metric. We first categorized the layouts into 5 difficulty intervals of $[150 + (i-1) \cdot 210, 150 + i \cdot 210]$, $i \in \{1, 2, 3, 4, 5\}$, representing from very simple to very difficult based on authors' playing experience. Then, we generated 2 layouts from each interval as the *reference layouts* and 3 layouts from each interval as the *evaluation layouts*. Due to the length limitation, we only show one layout for $i = 1, 3, 5$ from the *reference layouts* and the *evaluation layouts* in Fig. 13. Prior to the user study, we told participants the goal of one-shot-all-sink as well as the hitting interfaces of the billiard application, and let them play the game for 2 minutes to make sure all participants be familiar with the game. We started the user study by showing all *reference layouts* to the user as well as their designated difficulty levels. Then, they were asked to assign the difficulty level for all *evaluation layouts* arranged in a random order in a 5-point Likert scale manner. At each layout evaluation, we showed both the ball trajectory and the striking video to the user, and let them have three trial plays to reproduce the solution. After that, they gave the difficult level. Such an experiment design intended to prevent the user from solving the layout with a different solution. The user study was conducted on 16 subjects who were aged between 22 and 26, with 15 males and 1 female, and there were 14 having experiences in playing the billiard in real life and 2 had never played any billiard-related video games. We show the result in Fig. 14, and calculate the Spearman's rank correlation coefficient to get $\rho = 0.9191, p \approx 0.0$. Since $\rho > 0.7$ and $p \approx 0.0 < 10^{-3}$, we can conclude that there exists strong positive correlation between the user assigned level and the designated level, which indicates that the proposed difficulty metric is capable of revealing the player's opinions.




Fig. 13. This shows several examples of the *reference layouts* and the *evaluation layouts* for the difficulty study of the billiard. We enlarge the balls twice for better visualization.




Fig. 14. This shows the user assigned difficulty levels of billiard layouts, while the designated difficulty level is highlighted in red in the legend.

VII. CONCLUSION AND FUTURE WORK

This work proposes the L-MCTS as a general framework to produce spatio-temporal content under the user-specified metric, while experts are excluded from the workflow to achieve the goal of automatic generation and massive production. The framework involves the L-system to semantically model the elements and interactions of the content, makes use of the inverse embedding technique to improve the search efficiency, and effectively explores the content space with the MCTS sampling strategy. Our experiments show that the proposed method is capable of generating massive amount of desired content comparing to other baselines, and the generality is also demonstrated by applying the framework to 3 applications. However, there exists a few issues which are planned to be solved in the future. First, we assume that the branches are totally independent and do not affect each other, while such assumption limits the possible interactions as well as the content variation. For example, one single objective ball can be hit with the cue ball multiple times before pocketing, while such behavior cannot be represented by our system. One possible solution might apply graph-based modeling for interactions, but we would need to further develop the semantic L-system denotation and deduction for it. Another issue comes from the choice of the objective metric. This paper chooses all involved metrics to be positively related to the string length, while this assumption may not always be true in practice. For example, measuring the similarities between two chain reaction scenes is in general irrelevant with the string length. Therefore, it is intriguing to validate the effectiveness of the framework on more arbitrary metrics.

APPENDIX A COMPONENT PROPERTY IN CHAIN REACTION

Due to the fact that we simulate the chain reaction with the unity built-in physical engine, we summarize the fixed properties of each component for reproducibility.

All the balls involved in the application are identical to have the radius being $0.075m$ and the mass being $10kg$. The dimensionality of each single domino is $0.18m \times 0.04m \times 0.3m$ and the mass is $5kg$. The radius of the cup mouse is $0.2m$, and its mass is $2kg$, which is same with the baffle used in the pulley set.

APPENDIX B SPEARMAN'S RANK CORRELATION COEFFICIENT

The Spearman's rank correlation coefficient between two variables, X and Y , is calculated as

$$\rho = \frac{\sum_i^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i^n (x_i - \bar{x})^2 \sum_i^n (y_i - \bar{y})^2}}, \quad (5)$$

where the x_i and y_i are the rank of the raw data X_i and Y_i , and n is the sample size.

REFERENCES

- [1] J. Talton, Y. Lou, S. Lesser, J. Duke, R. Měch, and V. Koltun, "Metropolis procedural modeling," *ACM Trans. Graph.*, vol. 30, no. 2, pp. 11:1–11:14, 2011.
- [2] O. Stava, S. Pirk, J. Kratt, B. Chen, R. Mundefinedch, O. Deussen, and B. Benes, "Inverse procedural modelling of trees," *Comput. Graph. Forum*, vol. 33, no. 6, p. 118–131, 2014.
- [3] S. Snodgrass and S. Ontañón, "A hierarchical approach to generating maps using markov chains," in *Proceedings of the Tenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, ser. AIIDE'14. AAAI Press, 2014, p. 59–65.
- [4] L. Ferreira and C. Toledo, "A search-based approach for generating angry birds levels," in *Proceedings of the 2014 IEEE Conference on Computational Intelligence and Games*, ser. CIG 2014. New York, NY, USA: IEEE, 2014, pp. 1–8.
- [5] M. Fridenfalk, "Algorithmic music composition for computer games based on l-system," in *Proceedings of the 2015 IEEE Games Entertainment Media Conference (GEM)*, ser. GEM 2015. New York, NY, USA: IEEE, 2015, pp. 1–6.
- [6] S. Summerville, S. Philip, and M. Mateas, "Mcmcts pcg 4 smb: Monte carlo tree search to guide platformer level generation," *Artificial Intelligence in the Game Design Process*, vol. 11, no. 1, pp. 68–74, 2015.
- [7] R. Roussel, M.-P. Cani, J.-C. Léon, and N. Mitra, "Designing chain reaction contraptions from causal graphs," *ACM Trans. Graph.*, vol. 38, no. 4, pp. 43:1–43:14, jul 2019.
- [8] S. Cheney and D. Forsyth, "Sampling plausible solutions to multi-body constraint problems," in *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, ser. SIGGRAPH '00. New York City, NY, USA: ACM Press/Addison-Wesley Publishing Co., 2000, pp. 219–228.
- [9] C. Twigg and D. James, "Many-worlds browsing for control of multi-body dynamics," *ACM Trans. Graph.*, vol. 26, no. 3, p. 14–es, 2007.
- [10] —, "Backward steps in rigid body simulation," in *ACM SIGGRAPH 2008 Papers*, ser. SIGGRAPH '08. New York, NY, USA: ACM, 2008.
- [11] P. Prusinkiewicz and A. Lindenmayer, *The Algorithmic Beauty of Plants*. Heidelberg, Berlin, Germany: Springer-Verlag, 1990.
- [12] P. Prusinkiewicz, M. James, and R. Měch, "Synthetic topiary," in *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '94. New York, NY, USA: ACM, 1994, pp. 351–358.
- [13] P. Prusinkiewicz, L. Mündermann, R. Karwowski, and B. Lane, "The use of positional information in the modeling of plants," in *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '01. New York, NY, USA: ACM, 2001, pp. 289–300.
- [14] Y. Parish and P. Müller, "Procedural modeling of cities," in *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '01. New York, NY, USA: ACM, 2001, pp. 301–308.
- [15] J.-E. Marvie, J. Perret, and K. Bouatouch, "The fl-system: A functional l-system for procedural geometric modeling," *The Vis. Comput.*, vol. 21, no. 5, pp. 329–339, 2005.
- [16] J. Popović, S. Seitz, M. Erdmann, Z. Popović, and A. Witkin, "Interactive manipulation of rigid body simulations," in *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '00. USA: ACM Press/Addison-Wesley Publishing Co., 2000, p. 209–217.
- [17] M. Anderson, E. McDaniel, and S. Cheney, "Constrained animation of flocks," in *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ser. SCA '03. Goslar, DEU: Eurographics Association, 2003, p. 286–297.
- [18] M. Gleicher, "Motion editing with spacetime constraints," in *Proceedings of the 1997 symposium on Interactive 3D graphics*, ser. I3D '97. New York, NY, USA: ACM, 1997, pp. 139–ff.
- [19] Y. Ye and C. Liu, "Synthesis of detailed hand manipulations using contact sampling," *ACM Trans. Graph.*, vol. 31, no. 4, 2012.
- [20] L. Liu, K. Yin, and B. Guo, "Improving sampling-based motion control," *Computer Graphics Forum*, vol. 34, no. 2, pp. 415–423, 2015.
- [21] R. Kondo, T. Kanai, and K.-i. Anjyo, "Directable animation of elastic objects," in *Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ser. SCA '05. New York, NY, USA: ACM, 2005, p. 127–134.
- [22] J. Barbič, M. da Silva, and J. Popović, "Deformable object animation using reduced optimal control," *ACM Trans. Graph.*, vol. 28, no. 3, jul 2009.
- [23] A. Treuille, A. McNamara, Z. Popović, and J. Stam, "Keyframe control of smoke simulations," *ACM Trans. Graph.*, vol. 22, no. 3, p. 716–723, 2003.
- [24] Z. Pan, J. Huang, Y. Tong, C. Zheng, and H. Bao, "Interactive localized liquid motion editing," *ACM Trans. Graph.*, vol. 32, no. 6, nov 2013.
- [25] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton, "A survey of monte carlo tree search methods," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 4, no. 1, pp. 1–43, 2012.
- [26] F.-Y. Meng, J. Tang, W. Pedrycz, and Q.-X. An, "Optimal interaction priority calculation from hesitant fuzzy preference relations based on the monte carlo simulation method for the acceptable consistency and consensus," *IEEE Transactions on Cybernetics*, vol. 51, no. 12, pp. 5871–5882, 2020.
- [27] D. Silver, A. Huang, C. Maddison, A. Guez, L. Sifre, G. Driessche, J. S., I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, pp. 484–489, 2016.
- [28] N. Qi, Z. Fan, M. Huo, D. Du, and C. Zhao, "Fast trajectory generation and asteroid sequence selection in multispacecraft for multiasteroid exploration," *IEEE Transactions on Cybernetics*, vol. 52, no. 7, pp. 6071–6082, 2020.
- [29] H. Hong, M. Jiang, and G. G. Yen, "Improving performance insensitivity of large-scale multiobjective optimization via monte carlo tree search," *IEEE Transactions on Cybernetics*, 2023.
- [30] L. Kocsis and C. Szepesvári, "Bandit based monte-carlo planning," in *Machine Learning: ECML 2006*. Berlin, Heidelberg, Germany: Springer Berlin Heidelberg, 2006, pp. 282–293.
- [31] D. G. Alciatore, *The illustrated principles of pool and billiards*. Dr. Dave Billiards Resources, 2004.
- [32] R. Byrne, *Byrne's Treasury of Trick Shots in Pool and Billiards*. Houghton Mifflin Harcourt, 1983.
- [33] C. Spearman, "The proof and measurement of association between two things." 1961.
- [34] L. Leclezio, A. Jansen, V. H. Whittemore, and P. J. de Vries, "Pilot validation of the tuberous sclerosis-associated neuropsychiatric disorders (tand) checklist," *Pediatric Neurology*, vol. 52, no. 1, pp. 16–24, 2015.



Chia-Hsing Chiu is currently a Ph.D student at Department of Computer Science and Information Engineering, the National Taiwan University of Science and Technology, Taiwan. His research interests include computer graphics and multimedia.



Wen-Kai Tai received his M.S. and Ph.D. degrees in Computer Science from National Chiao Tung University in 1989 and 1995, respectively. He is currently a professor in the Department of Computer Science and Information Engineering at National Taiwan University of Science and Technology, Taiwan. His research interests include animation, real-time rendering, and gaming.



Yu-Chi Lai is currently a professor in Computer Science and Information Engineering with the National Taiwan University of Science and Technology, Taiwan, and his research interests include computer graphics and multimedia. He received the M.S. and Ph.D. degrees in CS from UW-Madison, Madison, WI, USA, in 2004 and 2010, respectively.



Andrew Chen is currently pursuing his graduate degree in Computer Science at the University of Illinois Urbana-Champaign, IL, USA. He completed his undergraduate studies at National Taiwan University of Science and Technology, Taiwan, where he earned dual bachelor degrees in both Computer Science and Information Engineering & Electrical Engineering. His research interests primary lie in computer graphics and human-computer interaction, with a focus on creativity-supporting design tools.



Shih-Syun Lin (Member, IEEE) received the Ph.D. degree in computer science and information engineering from National Cheng-Kung University (NCKU), Tainan City, Taiwan, in 2015. Since 2016, he has been a Faculty Member with the Department of Computer Science and Engineering, National Taiwan Ocean University (NTOU), Keelung City, Taiwan, where he is currently an Associate Professor with the Department of Computer Science and Engineering. He also leads the Intelligent Graphics Laboratory. His research interests include computer graphics, information visualization, information platform, and pattern recognition.



Pai-Yi Su is currently a M.A student at Department of Computer Science and Engineering, the National Taiwan Ocean University, Taiwan. His research interests include computer graphics and image processing.



Wei-Yao Ku completed a master's in 2022, and is currently a Ph.D. student in Computer Science and Engineering of the National Taiwan University of Science and Technology, Taiwan. His research interests involve computer graphics and physical simulation.