Infinite Art Gallery: A Game World of Interactively Evolved Artwork

Bryan Hollingsworth[1] and Jacob Schrum[2]

Abstract— **Procedural Content Generation (PCG) has been used extensively in video games as both a cost saving measure and a means to increase replayability. Evolutionary computation is an approach to PCG that has the ability to continually create new content tailored to particular users. This paper presents the Infinite Art Gallery, a game that uses established methods of evolving art with Compositional Pattern Producing Networks to allow users to explore a world of art tailored to their preferences. From a first-person perspective, users explore room after room of evolved paintings and sculptures. Each painting is a doorway that leads to a new room with art derived from a selected painting and sculpture in the previous room. Users have an inventory of collected paintings that can be used to return to old favorites, and they can influence the generation process with special item pickups. To gauge user response to the game, a human subject study was conducted with 30 users. Average overall enjoyment was 4.2 on a 5 point scale, indicating a general appreciation of the game. However, users also identified several areas in which the game could be improved, which will be pursued in future work.**

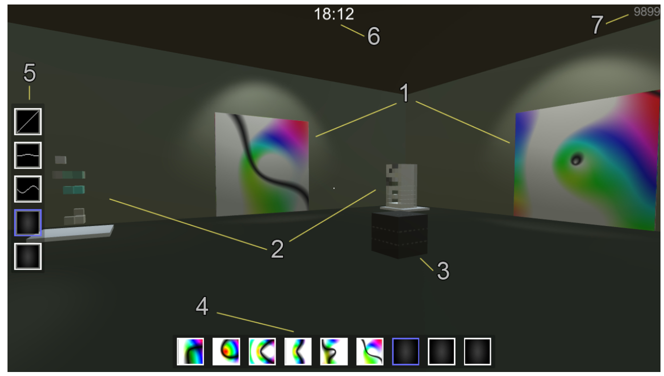

Fig. 1. **Infinite Art Gallery.** Users explore the art and walk through paintings to generate new art. Important game elements are labelled. 1) 2D paintings on walls. 2) 3D sculptures in corners. 3) Activation function cube pickup spawned in center of room. 4) 2D artwork inventory bar. 5) Collected activation function bar. 6) Time remaining. 7) Test subject ID. The time remaining and subject ID are only present in the human subject study.

I. INTRODUCTION

Procedural Content Generation (PCG) [1] can be used to circumvent the need for a large team of designers to spend a lot of time creating every component of a game. PCG also increases the replayability of a game, because there is always a chance of encountering something genuinely new and different on every playthrough. In some games, the entire focus is on generating new content, which is a novel type of open-ended experience for players.

This paper presents the Infinite Art Gallery (IAG; Fig. 1), a first-person exploration game in which interactive evolution continually produces new paintings and sculptures for the player to enjoy. The 2D paintings are generated as in the interactive art evolution system Picbreeder [2], and 3D sculptures are generated as in Endless Forms [3], both of which rely on Compositional Pattern Producing Networks (CPPNs [4]), a type of neural network with arbitrary activation functions (Section II-B).

The game lets users see how interesting their generated content can become. Every room the user enters has 2D paintings on the walls and 3D sculptures in the corners. The user can select one sculpture and walk through one painting. Each painting is a portal that leads to another room with newly evolved art derived from both the selected portal and sculpture. The art starts simple, but becomes more complex, while also being tailored to the specific player's preferences. The goal of the game is to create an enjoyable and engaging open-ended experience for users.

User enjoyment of the game was assessed via a human subject study with 30 participants. Users played the game for a set duration and then completed a short survey, in which they rated various aspects of the game and provided qualitative feedback. The results indicate that users enjoyed the game, particularly the 2D artwork, and felt as though their in-game decisions had a strong influence on the generated art. However, users were less impressed with the 3D sculptures, and had valuable suggestions for future improvements. The source code as well as an executable build of the version used in the human subject study are available at `https://southwestern.edu/~schrum2/SCOPE/iag.php`. This site also has selected art and videos from the study.

II. PREVIOUS WORK

This paper combines work in two major areas: PCG in games, and evolved CPPN art.

A. Procedural Content Generation

Procedural Content Generation (PCG) is the "algorithmic creation of game content with limited or indirect user input" [5]. PCG has been applied in many commercial games, such as *Rogue*, *Diablo*, *Spore*, and *Borderlands*. PCG can be applied to nearly any game element, such as levels, maps, characters, weapons, vehicles, textures, and sounds.

One means of PCG that is seldom utilized in commercial games, but whose promise has been explored in academic literature, is evolutionary computation. This approach, inspired

[1]B. Hollingsworth is attending Southwestern University, Georgetown, TX 78626, USA `hollingb@southwestern.edu`

[2]J. Schrum is an Assistant Professor of Computer Science at Southwestern University, Georgetown, TX 78626, USA `schrum2@southwestern.edu`

by Darwinian natural selection, involves selectively favoring and reproducing, with slight variation, content that players favor either directly or indirectly. *Galactic Arms Race* [6] is a space shooter in which the firing patterns of weapons evolve based on usage frequency. Researchers have also evolved levels for *Super Mario Bros* [7], various aspects of a tower defense game [8], and maps for a First-Person Shooter [9].

In these applications, evolved content is a way of enhancing a game in which the player(s) already have some goal. However, there are also games in which the purpose of the game is to explore evolved content. One example is *Petalz* [10], a social game in which users cultivate aesthetically pleasing flowers, which are generated by evolved Compositional Pattern Producing Networks (CPPNs [4]). Another game that uses CPPNs is *Artefacts* [11], in which users evolve 3D shapes to build interesting objects with, in a sandbox setting similar to *Minecraft*. This style of game, in which free-form exploration of evolved content is the main focus, is most similar to the Infinite Art Gallery (IAG) presented in this paper. Because IAG also depends on evolved CPPNs, the use of evolved CPPNs is discussed next.

B. Art via Compositional Pattern Producing Networks

Evolutionary computation has long been used to alter or generate art [12], [13], [14], [15]. The art generation approach in this paper extends work using CPPNs [4] to indirectly encode various forms of art, but any evolutionary art technique could be applied. CPPNs are a way to compose various functions in interesting ways, and are represented by artificial neural networks with arbitrary activation functions in their neurons. The set of available functions typically includes symmetrical, repeating, and asymmetrical patterns. Additionally, CPPNs are queried across a coordinate frame, such as x/y coordinates in a picture. The resulting artifact has a resolution determined by the number of queries across the coordinate frame, so the CPPN is capable of producing artifacts of arbitrary resolution.

These CPPNs are evolved using NeuroEvolution of Augmenting Topologies (NEAT [16]), an algorithm for evolving artificial neural networks. NEAT starts with simple networks lacking hidden neurons, whose topologies and weights are adjusted through mutations, such as splicing new nodes along existing links, altering link weights, and adding new links between existing nodes. When evolving CPPNs, a mutation to alter activation functions is sometimes also used, as in this paper. NEAT also includes a way of aligning networks for crossover, though this feature is not used in this paper.

An early demonstration of the artistic power of CPPNs was Picbreeder [2], an online interactive evolution system that draws pictures using CPPNs. The CPPNs map pixel locations to colors, producing an artistic image when repeated for all pixels. Users are presented with art generated by a population of CPPNs, and select art that they like to proceed to the next generation. Picbreeder also allows users to save their creations for others to view or use as starting points for evolving their own art. Pictures in IAG are generated using the same approach as Picbreeder, detailed in Section III-A.

A followup to Picbreeder called Endless Forms [3] extends CPPNs to evolve 3D shapes. Endless Forms is also an interactive evolution system, and functions similarly to Picbreeder, except that the input coordinates for the CPPNs are voxel locations in three dimensions, and instead of mapping to colors, the CPPNs determine whether or not a given voxel is present in the 3D shape. Shapes are rendered in a single color of the user's choice. In order to prevent the final results from looking blocky, the Marching Cubes algorithm [17] smooths the edges of the evolved shapes.

This same encoding has been extended to allow individual voxels to have distinct colors [18]. A version called 3DObjectBreeder [19] assigns colors to voxels, but does not smooth the shapes. This purposefully blocky output is reminiscent of the artistic style of *Minecraft*. This unsmoothed approach is used in IAG, and described in Section III-B.

In the same work by Tweraser et al. [19], CPPNs are queried across time to evolve 2D and 3D animations. Others have used CPPNs to produce audio timbres that when sampled at different frequencies correspond to different notes, so that a MIDI file can use the CPPN as an instrument to play a song [20]. CPPN controllers have also produced art-generating agents in an open-ended artificial life domain [21], and there are several examples of CPPNs generating soft-bodied robots [22], [23], [24]. In addition to being artistically appealing, such CPPN-agents could be adapted to serve as enemies in a video game setting. These applications of CPPNs are not incorporated into the current version of IAG, but are all intriguing possibilities for future work.

III. Infinite Art Gallery

The Infinite Art Gallery (IAG) was developed using the *Unity* engine, with backend code written in C#. This section describes how IAG generates evolutionary art, resulting in an open-ended exploration game.

IAG inserts a player into a single square room decorated with artwork created by CPPNs (Fig. 1). The room contains both 2D *paintings* and 3D *sculptures*. The player sees the room from a first-person perspective, and is free to move around and examine the artwork. Each painting and sculpture is created by a CPPN. There is one painting on each wall and one sculpture in each corner. Each painting serves as a portal to move the player to the next room. Once a player has chosen a painting from the presented options, they walk through that painting, and IAG uses their selection to evolve three new paintings. If a sculpture was selected, then the new room will also contain four new evolved sculptures.

After a player has moved into a new room, the portal they just exited from is replaced with a solid white painting. This image represents a return portal the player can use to revisit the previous room. Travel in any direction that is not the return direction will always be into a new room. It is not possible in IAG to enter a previously generated room without backtracking. Using this system, players can move though the gallery for any amount of time, yet still retrace their steps all the way to the original room if desired.

A. 2D Paintings

IAG has different types of CPPNs for paintings and sculptures. At the start of the game, four random painting CPPNs are generated using the same input-output mapping as Picbreeder [2]. Each CPPN has four inputs and three outputs (Fig. 2). Initial networks are fully connected with no hidden neurons and only identity as an activation function. Each CPPN generates a painting on one wall of the starting room.

Each painting is generated by querying all x/y coordinates in an image with a CPPN. Paintings in IAG have a resolution of 128×128 pixels. For each pixel, the four CPPN inputs are the x coordinate, y coordinate, distance of the pixel from the center, and a constant bias of 1. The x and y coordinates are mapped to the range of $[-1, 1]$ to create symmetry across these axes. The distance from the center is calculated with these scaled coordinates, and then multiplied by $\sqrt{2}$, as in the original Picbreeder. Though not strictly necessary, the center distance makes radial patterns easy to generate.

These inputs pass though the links and neurons of the neural network in order to generate a color for the associated pixel based on the CPPN outputs. The three outputs represent the hue, saturation, and brightness value of the pixel (HSV). The network output ranges depend on the activation functions in the output neurons, which in some cases allow arbitrarily large or small values. However, HSV encoding requires values in specific ranges.

Hue must be in the range $[0, 1]$. Simply clamping all values to this range would tend to saturate the hue to be red. Instead, outputs are scaled to the range $[0, 1]$ with respect to the minimum and maximum outputs that the CPPN produces across all pixels that are queried. In contrast, saturation is simply clamped to the range $[0, 1]$, and brightness is clamped to the range $[-1, 1]$ before taking the absolute value in order to confine it to the range $[0, 1]$. This unusual design choice, consistent with Picbreeder, is desirable because it tends to produce dark lines that delineate shapes in the images.

B. 3D Sculptures

Each IAG room contains a sculpture in each corner. Each sculpture is generated by its own CPPN in a manner similar to the paintings, but more inputs and outputs are required in these CPPNs (Fig. 3). This encoding is inspired by Endless Forms [3], though adheres more closely to the more recent 3DObjectBreeder [19]. As with paintings, the initial CPPNs are fully connected, but have no hidden neurons.

Instead of querying each pixel in a 2D space, each voxel in a 3D space is queried. Each sculpture is made up of $5 \times 5 \times 10$ voxels, which are cubes measuring 0.5 Unity units on each side. Therefore, a z coordinate input is needed in addition to x and y. There is still a distance from center input, but now there are also inputs for distances from the center within the xy, xz, and yz planes. The constant bias of 1.0 is also present. As with paintings, coordinates are mapped to the range $[-1, 1]$, and distances are calculated with respect to this modified range before being multiplied by $\sqrt{2}$.

The HSV outputs are mapped in the same manner as in paintings, and determine the color of the given voxel, but only if it is present in the sculpture. A new output determines whether the given voxel is even present. The voxel is present if the CPPN output is greater than a threshold of 0.1.

A new feature in this paper is that this additional output can also serve as an alpha value for the voxel. When each sculpture is created, transparency is randomly enabled or disabled. If transparency is disabled, then this output behaves the same as in Endless Forms [3] and 3DObjectBreeder [19]. However, when enabled, a CPPN output above the presence threshold instead becomes the alpha value of the voxel. Therefore, alpha values below 0.1 are not possible.

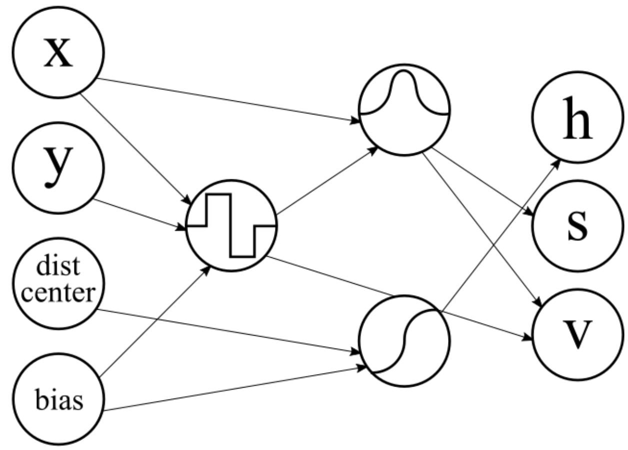

Fig. 2. **Example CPPN For Generating 2D Paintings.** Structure of a CPPN with four input nodes, three hidden layer nodes, and three output nodes. 2D artwork is created by querying every x/y coordinate in the image along with the distance of that point from the center of the image and a constant bias value. Inputs are transmitted though links and nodes containing different activation functions. Each link has a weight that modifies the signal via multiplication as it travels through the network. The final outputs encode the hue, saturation, and brightness value of the given pixel.

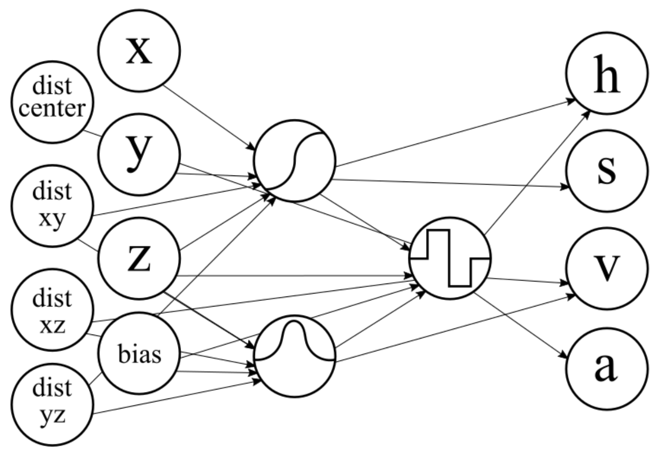

Fig. 3. **Example CPPN For Generating 3D Sculptures.** Structure of a CPPN with eight input nodes, three hidden layer nodes, and four output nodes. 3D artwork is created similarly to 2D artwork with the additional inputs of z, and distances from center in the xy, xz, and yz planes. The additional alpha output is used to determine the presence of a voxel in the final artifact, or if transparency is enabled for that sculpture, to set the alpha value of the voxel.

A CPPN output of 0.5 would represent 50% transparency, and an output of 1.0 would be completely solid. Any output above 1.0 is restricted to be 1.0. The use of transparency is effective in the setting of IAG because there is actually an environment that can be seen behind the sculptures.

C. Mutations

With both 2D and 3D artwork, evolution occurs through mutations. Once an individual is selected and evolution is triggered, multiple offspring are created and mutated. For both 2D and 3D artwork, one of the offspring is mutated only once and placed in the same location in the new room. In the case of 2D artwork, the paintings to the left and right of the selected individual are mutated two to six times, determined randomly. Since the entry portal becomes a return portal, there are only three new paintings in each new room. In the case of 3D sculptures, the three remaining sculpture offspring are mutated two to six times, determined randomly.

A CPPN can be mutated by adding a new link between existing nodes, splicing a new node with a random activation function into an existing link, altering the weights of links in the CPPN, or changing the activation function of an existing node. Whenever a mutation action occurs, exactly one of these four options is chosen with equal probability. It is important to assure frequent and significant mutations so that there will actually be variation from room to room.

Adding a new link can connect a neuron to any neuron that occurs later in the network, meaning closer to the outputs. Recurrent links are disabled. Splicing adds a new node along an existing link and sets the link weights on each side randomly. The new node is given a random activation function. Link perturbation gives each link a chance to be mutated with a random probability that is itself randomly generated whenever link perturbation occurs. Therefore, the threshold for individual links to be mutated could randomly be low or high for any given mutation, which allows for sudden large changes. Each link weight that gets mutated is modified via Gaussian perturbation. Finally, a randomly selected node can change to a random activation function. Any time a random activation function is needed, it is selected from functions the player has collected in IAG (Section III-F).

D. Interaction With Sculptures

Interaction with sculptures is different than with paintings. There are two actions associated with sculptures: selection for mutation and random replacement.

Left-clicking on a sculpture selects that sculpture, indicated by the base turning green. Only one sculpture can be selected for mutation. If a sculpture is selected when the player leaves the room through a 2D portal, then all four sculptures in the next room will be mutated offspring of the selected sculpture from the previous room. New sculptures are generated even if using the white return portal, if a sculpture was selected. Similarly, if no sculpture is selected, then all sculptures in the room will remain the same whether using return portals or art portals.

Right-clicking on a sculpture will replace it with a random sculpture. A new CPPN is generated using the player's collected activation functions (Section III-F), and a replacement sculpture appears. The new sculpture CPPN will be fully connected, but have no hidden neurons, as if it were a sculpture in the initial room.

E. Painting Inventory

Players can left click on a painting to save a copy in an inventory bar at the bottom of the screen. Saving a painting populates the inventory slot with a thumbnail image of the artwork. Up to nine paintings can be saved. Players can use a saved painting to replace a wall painting in a room by selecting a painting in the inventory with the mouse wheel and right-clicking on the wall painting they want to replace. Once a painting has been used it is removed from the inventory. This feature allows a player to store an appealing painting, and wait until later to explore its offspring. This feature is particularly useful if the user wants to focus exclusively on evolving sculptures for a while, and return to an appealing painting later.

Sculptures cannot be saved in the inventory, though this or a similar mechanic may be included in future releases.

F. Collecting Activation Functions

At the start of the game, the only activation function used to generate art is identity. This function is represented in a vertical bar on the left side of the screen. However, rooms have a 33% chance to spawn item pickups in the center that represent new activation functions for the CPPNs to use. The items are represented as black cubes that show the characteristic graph of the function in white. Once a function has been picked up, it is added to the vertical inventory bar on the left. This inventory contains five spaces for activation functions. Users can dispose of functions at any time, but must always keep at least one. When picking up new functions, users can have them replace any currently possessed function rather than simply adding it to the set.

The activation function pickups give the player more control over the evolution of the networks. Features that emerge in images can be very dependent on what functions are available. For instance, the square and sawtooth waves produce distinct banding and repetition in the artwork. Sine functions also produce repetition, but with smoother edges.

The complete list of activation functions that can be found includes identity, tanh, square wave, Gaussian, sine, sawtooth wave, and absolute value. Some of these functions were used in Picbreeder [2] and Endless Forms [3], though several were not. However, work by Tweraser et al. [19] demonstrated that these additional functions can result in qualitatively distinct and interesting patterns in generated art.

IV. HUMAN SUBJECT STUDY

A human subject study was conducted to evaluate the effectiveness of IAG in generating interesting and meaningful content as well as an engaging experience for players. A total of 30 participants were given 20 minutes to explore the

TABLE I

User Enjoyment And Desire to Play More

User ratings of their enjoyment of different components (percentages and exact numbers). **Desire** indicates the user's desire to play more of the game.

	Less				More
Rating	1	2	3	4	5
Game	0%(0)	0%(0)	$16.\overline{6}\%$(5)	$46.\overline{6}\%$(14)	$36.\overline{6}\%$(11)
2D Art	0%(0)	0%(0)	$13.\overline{3}\%$(4)	$33.\overline{3}\%$(10)	$53.\overline{3}\%$(16)
3D Art	$3.\overline{3}\%$(1)	$23.\overline{3}\%$(7)	$33.\overline{3}\%$(10)	$26.\overline{6}\%$(8)	$13.\overline{3}\%$(4)
Desire	$3.\overline{3}\%$(1)	$6.\overline{6}\%$(2)	20%(6)	$33.\overline{3}\%$(10)	$36.\overline{6}\%$(11)

gallery and given a questionnaire upon completion. In addition to the questionnaire, all 2D and 3D artwork generated by each user was saved, and each play session was recorded using screen capture software. Participants were students, faculty, and staff at Southwestern University in Georgetown, Texas, USA. The goal of the study was to collect feedback on the game mechanics and gauge player perception of their ability to influence and direct the evolution of the artwork.

A. *Procedure*

Subjects were given a brief explanation on how to interact with IAG but provided very few details on how art was being generated. They were told how to select 2D paintings and 3D sculptures that they liked and that those selections would be used to generate new artwork, but they were not told exactly how this would happen. Subjects were told that the black function pickups were related to potential features in the evolved art, but again, were not given any specifics as to how those features would or could manifest.

This information was given before IAG was launched and was scripted to ensure uniform player knowledge for each study. Once IAG was launched, the major features of the game, such as movement and inventory controls, were repeated as instructional splash screens to ensure that all information from the tutorial was understood and there were no questions. Only after the information screens had been dismissed was the timer started. At this point the investigator left the room. While each subject was given the opportunity to pause the experiment to ask the investigator questions, none of the subjects did so. The executable for the version of the game that the participants played is available for download at `https://southwestern.edu/~schrum2/SCOPE/iag.php`.

B. *Quantitative Results*

Subjects were asked to answer a survey following their experience in IAG. The survey consisted of 15 questions with 11 questions asking the subjects to rate specific topics on a scale of 1 to 5. Table I reports the level of enjoyment of the game and both types of art (2D and 3D), as well as the player's desire to play the game again. Most subjects enjoyed the experience provided by IAG, as indicated by an average enjoyment score of 4.2. More than half ($53.\overline{3}\%$) of the users rated their enjoyment of the 2D artwork at a 5, with an average rating of 4.4. Responses to the 3D artwork were less favorable, with an average rating of $3.2\overline{3}$, including eight scores below 3. Most subjects were interested in playing the game again, as the average score associated with this question was $3.9\overline{3}$, although three individuals did give scores less than 3 to this question.

TABLE II

Ratings of Individual Game Components

User ratings of game components pertaining to both 2D and 3D art (percentages and exact numbers). **Variation** is the diversity of art of the given type seen by the user. **Choice** measures the impact of the user's choices on what art was generated. **Items** measures the impact of the activation function pickups on the generated art. **Inventory** measures usefulness of the inventory bar at the bottom of the screen (2D art only). Note that one user did not answer the question about variation in 3D art, leading to unusual and approximate percentage calculations.

	Less				More
Rating	1	2	3	4	5
			2D Art		
Variation	0%(0)	10%(3)	$16.\overline{6}\%$(5)	50%(15)	$23.\overline{3}\%$(7)
Choice	0%(0)	$3.\overline{3}\%$(1)	$6.\overline{6}\%$(2)	40%(12)	50%(15)
Items	10%(3)	20%(6)	30%(9)	$26.\overline{6}\%$(8)	$13.\overline{3}\%$(4)
Inventory	$3.\overline{3}\%$(1)	$3.\overline{3}\%$(1)	20%(6)	$16.\overline{6}\%$(5)	$56.\overline{6}\%$(17)
			3D Art		
Variation	0%(0)	34.5%(10)	31.0%(9)	20.7%(6)	13.8%(4)
Choice	$6.\overline{6}\%$(2)	$33.\overline{3}\%$(10)	$23.\overline{3}\%$(7)	20%(6)	$16.\overline{6}\%$(5)
Items	$43.\overline{3}\%$(13)	20%(6)	30%(9)	$6.\overline{6}\%$(2)	0%(0)

Table II reports the distribution of subject responses for specific questions concerning the 2D and 3D artwork. Subjects responded more positively to the 2D paintings than the 3D sculptures with regards to variation, control, and function cube impact. Most subjects felt that the variation in the 2D artwork was high, with an average rating of $3.8\overline{6}$. Fewer individuals perceived variation in the 3D artwork, with most of the results being 2 and 3, producing an average rating of approximately 3.14. Subjects were asked to rate the amount of control they felt they had in influencing the generated artwork by their choices. Their ratings were high for 2D paintings (average $4.3\overline{6}$), with most giving a ranking above 3, but low for 3D sculptures (average $3.0\overline{6}$). Fig. 4 shows how the art witnessed by three different users changed as they played IAG, and demonstrates how much variation there is both within individual sessions and across sessions. The impact of activation function pickups was not perceived to be high, with a middling average score of $3.1\overline{3}$ for 2D paintings and a dismal average of 2 for 3D sculptures, which lacks any ratings of 5. For reference, an example of how function pickups affected the art seen by one user is shown in Fig. 5. Finally, the use of the inventory to save and replace 2D artwork was viewed positively with an average rating of 4.2, the majority of which were a 5 ($56.\overline{6}\%$). The 3D artwork did not have an inventory, so there was no ability to save a sculpture to be used later. The differences in the 2D and 3D ratings are discussed in more detail in Section V.

These results indicate that IAG can deliver a meaningful and enjoyable experience to players while suggesting areas for improvement concerning the collection of functions and with the presentation of the 3D artwork. Users enjoyed the game mechanics around the 2D artwork, including the use of

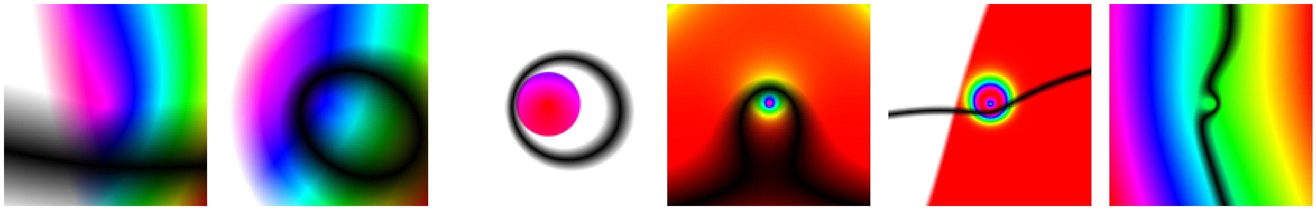

(a) Subject A, Image 1 (b) Subject A, Image 10 (c) Subject A, Image 17 (d) Subject A, Image 62 (e) Subject A, Image 134 (f) Subject A, Image 144

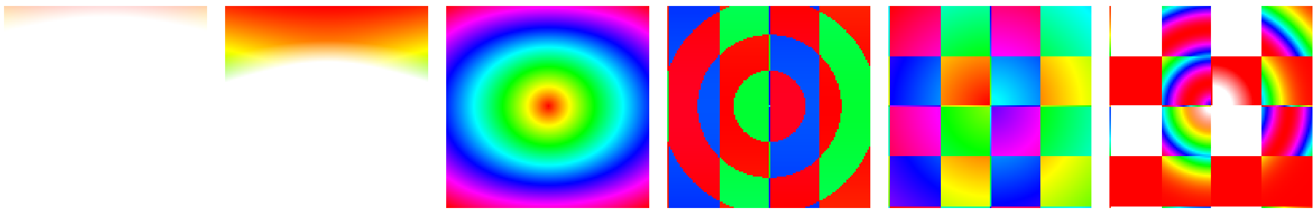

(g) Subject B, Image 3 (h) Subject B, Image 11 (i) Subject B, Image 27 (j) Subject B, Image 32 (k) Subject B, Image 63 (l) Subject B, Image 94

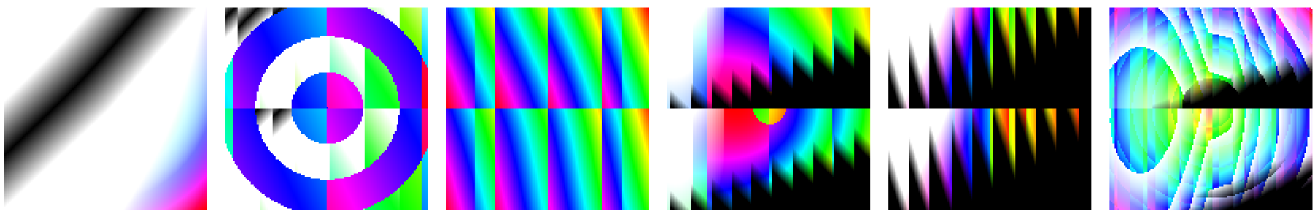

(m) Subject C, Image 1 (n) Subject C, Image 10 (o) Subject C, Image 18 (p) Subject C, Image 25 (q) Subject C, Image 31 (r) Subject C, Image 35

Fig. 4. **Comparison of 2D Paintings Selected by Three Test Subjects.** Six selected images each from three human subjects are shown to illustrate the variation in play sessions. Each painting was encountered by the subject and subsequently selected, with paintings to the left encountered earlier and paintings to the right encountered later. Many subjects started with images similar to the starting images shown here, but through the course of selective breeding, each user ended up with different outcomes.

the inventory. Insight into the lower ratings for 3D artwork comes from the short answer questions.

C. *Qualitative Results*

Subjects were asked to provide short answers to four questions. They were asked to describe a specific 2D painting and 3D sculpture that they found most interesting, and were asked what qualities or features of the game would make them want and not want to play again. Despite the 2D artwork receiving higher ratings overall, most of the descriptions provided were abstract, focusing on shapes and colors. One response said the artwork looked like "tie dye," and another said, "It was like a Jackson Pollock painting." Only four users said they saw familiar imagery in the paintings: "HAL 9000 from 2001: A Space Odyssey," and "Earth's core" being the most specific. One user described "A face with one eye, triangular shaped" and another described "strange outer space pictures of truly weird planets." Many users responded that the colors and contrasts provided were enjoyable.

Responses for the 3D artwork contained many more comparisons to recognizable objects, the majority being living creatures. Users described sculptures as looking like birds more than any other animal. One user described a "blocky statue of a bird," where two others were more specific referring to one as a "flamingo" and another as a "blue and green owl!" Other animal forms included a "pixelated version of a rabbit from Wallace and Gromit" and a "small man with a orange head and a black body." Sculptures with a lot of negative space appealed to some users, providing descriptions of levitating shapes or concave aspects. Translucent blocks appearing in sculptures was mentioned in seven of the responses. As with the 2D artwork, color and contrast were features that users found enjoyable.

When asked what features would make the user want to play the game more, the most common theme in the responses was control. Some users showed an interest learning how to manipulate the artwork that was generated, with one user describing a specific event they would like to test involving the mixing of colors in different artwork. The ability to keep the generated artwork was requested by more than one user, with one wanting to use the game to create art for their home. Users described the experience as relaxing and they enjoyed the atmosphere and freedom the game provided. Use of the inventory was mentioned by multiple users as a positive feature that kept the game play enjoyable. Some users enjoyed the selection mechanics of walking through a painting to evolve the artwork.

When asked about features that would make the user not want to play more, many indicated that changes in the artwork were too subtle during evolution. Responses indicated frustration with colors or shapes becoming repetitive. The 3D artwork was not as interesting or controllable as the 2D artwork. Some users reported eventually ignoring the sculptures and only focusing on the paintings. The use of the function cubes was confusing to some users. Responses indicated trouble determining how and when they were used

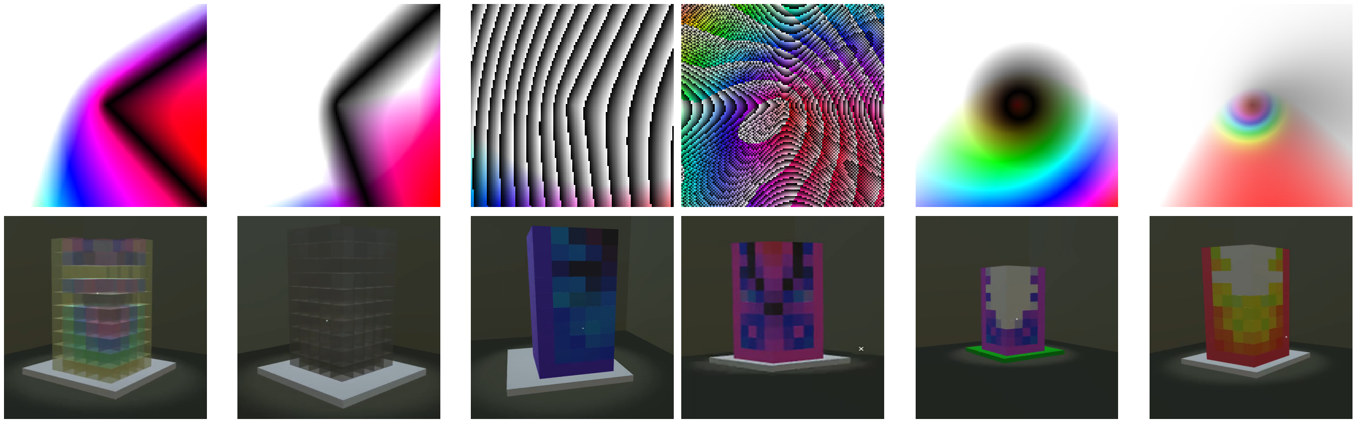

Fig. 5. **Artwork Evolution After Function Pickups.** Each column shows a 2D painting and 3D sculpture generated after the designated activation function pickup was grabbed. These figures show how activation function pickups can change the art that is generated, though users generally did not appreciate the impact of these items. First absolute value was obtained, but it was then replaced with the sawtooth wave function. Then tanh and Gaussian were added. Next the user loaded the pictured painting from their inventory. The sculpture shown in this column was not affected by this action, and is simply a result that coincided with the selection of that painting. Finally, the sine function pickup was obtained. There seems to be significant variation in the 2D paintings, but users did not attribute these changes to the activation functions. The 3D sculptures show variation in their color patterns, but admittedly show little variance in shape.

or what reason there was to collect them or to manage their collections. In terms of gameplay and mechanics, the lack of a specific goal or objective was viewed negatively by a few users, with one user mentioning a lack of reward mechanisms. Some users found the controls to be confusing regarding either the saving or use of an inventory item, or the way function cubes replaced a selected cube in the collection. One subject did not provide an answer to this question.

V. Discussion and Future Work

The responses of the human subjects show that interaction with an evolutionary algorithm in a video game environment can provide an enjoyable experience with a high desirability to play more. The use of game mechanics to interact with the evolution of CPPNs had mixed results, with a clear difference between 2D and 3D artwork. Subject perception of their ability to influence the artwork being generated was high for 2D artwork and low for 3D artwork. The mechanism for mutation on both types of artwork was the same, but 3D sculptures may have been harder to influence, both because their CPPNs involve more inputs, and because changes in the presence of voxels only occur when a threshold is crossed rather than varying smoothly across a continuum. The perceived impact of the activation functions on evolution was generally low, though these responses seem to indicate confusion about how the activation functions are used. It is perhaps hard to separate the influence of activation functions from the impact of randomness in mutations, to the untrained eye. Use of the inventory for 2D artwork was seen as the most impactful control the subject had over the outcome of the experience, indicating that a way to curate 3D sculptures with a similar mechanism should be added.

Examination of the differences in experience provided by 2D and 3D artwork suggest a few areas to target for improvement. The perceived lack of variation in the 3D artwork could be addressed by increasing the resolution of the artwork (more voxels), or by implementing Marching Cubes [17] to smooth and round objects. Randomizing or allowing the user to control the threshold at which a CPPN output causes a voxel to be present could also lead to more diverse shapes. The limited width and depth of the 3D sculptures potentially made the application of various activation functions hard to perceive, resulting in the overwhelmingly negative perception of these functions. However, even with lower ratings concerning the 3D sculptures, subjects' descriptions showed a higher frequency of identifying sculptures as recognizable objects, which demonstrates a positive interplay between computational creativity and human imagination.

Subjects reported a high amount of variation in the 2D artwork. Examination of the generated artifacts shows a high amount of variation over all subjects, but also shows a tendency for certain features like color, patterns, and complexity to be dominant within an evolution session. This outcome is expected as the new artwork is generated based on the preferences of a single user, but subject perception was that the variation was still high. The use of the inventory to switch between paintings seems to play a large role in the subject's perception. The lack of an inventory for 3D sculptures likely influenced the low perception of variation, and suggests a need to add that feature. The sculpture inventory could be a separate inventory, or part of the existing one.

CPPN crossover was not used in the study, because users could only ever select one item of each type at a time. In the future, the inventory could provide partners for crossover. Specifically, selected artwork could be crossed with a randomly selected item in the inventory to generate new artwork for the new room. This feature would add more meaning to the inventory as a game mechanic.

Player perception of activation function pickups could be improved by more drastically demonstrating how activation functions influence the generated art. For instance, the ability to use a pickup to replace all activation functions in a given CPPN with the designated function could be added. Having an example of the patterns common to one particular activation function may make those patterns more appreciated when seen combined with others. Other items could control increasing or decreasing the number of mutations per evolution or changing the chances of certain mutations occurring. Optimizing the rate at which the items spawn and the magnitude of their effect is also a topic for future work.

The addition of other evolved artifacts could enhance and modify the user experience. For example, MIDI music can use an evolved CPPN as an instrument [20], providing another way for users to generate content that appeals to them. Users could have the option to play music from different radio stations that gradually mutate over time. Animating some of the CPPN artwork [19] could also create a more engaging experience. Another type of evolved CPPN artifact that could add an element of danger and conflict to the game would be soft-bodied robots [22], [24], [23], which could be rewarded for damaging the player. Users could be surprised by 3D sculptures suddenly and unexpectedly coming to life. Ultimately, the diversity of applications for CPPNs means that there are many potential enhancements to this game world that could take advantage of evolutionary computation. As the types of content are expanded, the potential feedback from users expands as well. While this phase of the project was focused on overall user enjoyment, the game could be assessed using other models in the future, such as the Hedonic-Motivation System Adoption Model [25].

VI. CONCLUSIONS

Procedural Content Generation has been used to create various components of games in the past and, as shown by a human subject study with the Infinite Art Gallery (IAG), can be used to create an enjoyable open-ended experience for players. This proof of concept shows that the use of generated artwork as the primary content in an exploration game has potential as a new type of gaming experience. Users liked the relaxing atmosphere as well as exploring and influencing the creation of new artwork. Feedback on the 2D artwork was strongly favorable, suggesting that the problems with 3D artwork can be overcome in the future. As noted, additional content can be generated and added to IAG, including music, animation, and soft robots, to enhance the experience and provide more options for players.

ACKNOWLEDGMENTS

The authors would like to thank the Schloss Dagstuhl team and the organizers of the Dagstuhl Seminar 17471 for a creative and productive seminar. This research is supported in part by the Summer Collaborative Opportunities and Experiences (SCOPE) program for undergraduates, funded by various donors to Southwestern University.

REFERENCES

[1] N. Shaker, J. Togelius, and M. J. Nelson, *Procedural Content Generation in Games*. Springer, 2016.

[2] J. Secretan, N. Beato, D. B. D'Ambrosio, A. Rodriguez, A. Campbell, J. T. Folsom-Kovarik, and K. O. Stanley, "Picbreeder: A Case Study in Collaborative Evolutionary Exploration of Design Space," *Evolutionary Computation*, vol. 19, no. 3, pp. 373–403, 2011.

[3] J. Clune and H. Lipson, "Evolving Three-dimensional Objects with a Generative Encoding Inspired by Developmental Biology," in *European Conference on Artificial Life*, 2011, pp. 141–148.

[4] K. O. Stanley, "Compositional Pattern Producing Networks: A Novel Abstraction of Development," *Genetic Programming and Evolvable Machines*, vol. 8, no. 2, pp. 131–162, 2007.

[5] J. Togelius, E. Kastbjerg, D. Schedl, and G. N. Yannakakis, "What is Procedural Content Generation? Mario on the Borderline," in *Workshop on Procedural Content Generation in Games*, 2011.

[6] E. J. Hastings, R. K. Guha, and K. O. Stanley, "Automatic Content Generation in the Galactic Arms Race Video Game," *IEEE Transactions on Computational Intelligence and AI in Games*, 2009.

[7] N. Shaker, M. Nicolau, G. N. Yannakakis, J. Togelius, and M. O'neill, "Evolving Levels for Super Mario Bros Using Grammatical Evolution," in *Computational Intelligence and Games*. IEEE, 2012.

[8] P. Avery, J. Togelius, E. Alistar, and R. P. van Leeuwen, "Computational Intelligence and Tower Defence Games," in *Congress on Evolutionary Computation*. IEEE, 2011, pp. 1084–1091.

[9] P. T. Olsted, B. Ma, and S. Risi, "Interactive Evolution of Levels for a Competitive Multiplayer FPS," in *Congress on Evolutionary Computation*. IEEE, 2015, pp. 1527–1534.

[10] S. Risi, J. Lehman, D. B. D'Ambrosio, R. Hall, and K. O. Stanley, "Petalz: Search-based Procedural Content Generation for the Casual Gamer," *IEEE Transactions on Computational Intelligence and AI in Games*, 2015.

[11] C. Patrascu and S. Risi, "Artefacts: Minecraft Meets Collaborative Interactive Evolution," in *Computational Intelligence and Games*. IEEE, 2016, pp. 349–356.

[12] S. Todd and W. Latham, *Evolutionary Art and Computers*. Orlando, FL, USA: Academic Press, Inc., 1994.

[13] G. Greenfield, "Robot Paintings Evolved Using Simulated Robots," in *Applications of Evolutionary Computing*, 2006, pp. 611–621.

[14] A. Neumann, Z. L. Szpak, W. Chojnacki, and F. Neumann, "Evolutionary Image Composition Using Feature Covariance Matrices," in *Genetic and Evolutionary Computation Conference*, 2017.

[15] C. Jacob, "Evolving Evolution Programs: Genetic Programming and L-systems," in *Genetic Programming*, 1996, pp. 107–115.

[16] K. O. Stanley and R. Miikkulainen, "Evolving Neural Networks Through Augmenting Topologies," *Evolutionary Computation*, 2002.

[17] W. E. Lorensen and H. E. Cline, "Marching Cubes: A High Resolution 3D Surface Construction Algorithm," in *Computer Graphics and Interactive Techniques*, 1987, pp. 163–169.

[18] J. Lehman, S. Risi, and J. Clune, "Creative Generation of 3D Objects with Deep Learning and Innovation Engines," in *International Conference on Computational Creativity*, 2016.

[19] I. Tweraser, L. E. Gillespie, and J. Schrum, "Querying Across Time to Interactively Evolve Animations," in *Genetic and Evolutionary Computation Conference*. ACM, 2018.

[20] B. T. Jónsson, A. K. Hoover, and S. Risi, "Interactively Evolving Compositional Sound Synthesis Networks," in *Genetic and Evolutionary Computation Conference*. ACM, 2015, pp. 321–328.

[21] K. O. Stanley, N. Cheney, and L. B. Soros, "How the Strictness of the Minimal Criterion Impacts Open-Ended Evolution," in *Artificial Life*, 2016, pp. 208–215.

[22] N. Cheney, R. MacCurdy, J. Clune, and H. Lipson, "Unshackling Evolution: Evolving Soft Robots with Multiple Materials and a Powerful Generative Encoding," in *Genetic and Evolutionary Computation Conference*. ACM, 2013.

[23] F. Corucci, N. Cheney, F. G. Serchi, J. Bongard, and C. Laschi, "Evolving Soft Locomotion in Aquatic and Terrestrial Environments: Effects of Material Properties and Environmental Transitions," *Soft Robotics*, 2018.

[24] S. Kriegman, N. Cheney, F. Corucci, and J. C. Bongard, "Interoceptive Robustness Through Environment-mediated Morphological Development," in *Genetic and Evolutionary Computation Conference*, 2018.

[25] P. Lowry, J. Gaskin, N. Twyman, B. Hammer, and T. Roberts, "Taking "Fun and Games" Seriously: Proposing the Hedonic-Motivation System Adoption Model (HMSAM)," *Journal of the Association of Information Systems*, vol. 14, no. 11, pp. 617–671, 2013.