SN Computer Science (2022) 3:505
https://doi.org/10.1007/s42979-022-01329-0

ORIGINAL RESEARCH

Axial Generation: Mixing Colour and Shapes to Automatically Form Diverse Digital Sculptures

Edward Easton[1] · **Anikó Ekárt**[1] · **Ulysses Bernardet**[1]

Received: 28 July 2021 / Accepted: 14 July 2022 / Published online: 5 October 2022

Abstract
Automated computer generation of aesthetically pleasing artwork has been the subject of research for several decades. The unsolved problem of interest is how to please any audience without requiring too much of their involvement in the process of creation. Two-dimensional pictures have received a lot of attention; however, 3D artwork has remained relatively unexplored. This paper showcases an extended version of the Axial Generation Process (AGP), a versatile generation algorithm that can create both 2D and 3D items within the Concretism art style. The extensions presented here include calculating colour values for the artwork, increasing the range of forms that can be created through dynamic sizing of shapes and including more primitive shape types, finally, 2D items can be created from multiple viewpoints. Both 2D and 3D items generated through the AGP were evaluated against a set of formal aesthetic measures and compared against two established generation systems, one based on manipulating pixels/voxels and another tracking the path of particles through 2D and 3D space. This initial evaluation shows that the process is capable of generating visually varied items which exhibit a generally diverse range of values across the measures used, in both two and three dimensions. Comparatively, against the established generation processes, the AGP shows a good balance of performance and ability to create complex and visually varied items.

Keywords Evolutionary computation · 2D and 3D art generation · Concretism

Introduction

The abstraction of real-world phenomena is a key element of art. The level of abstraction depends on many different factors, such as the subject of the artwork, the medium the artwork uses and the style it is created in. 3D media allows more accurate representations of real-world objects to be created than 2D media, which often abstracts by flattening 3D items onto a 2D canvas. The style of a piece of art often transcends across both 2D and 3D media and heavily influences the subject. A style can be thought of as the combination of aspects that can be used to identify and categorise pieces of art. These categories are used as a method of identifying and comparing both artists and their work. For paintings and sculptures, these styles are often well-defined, Van Gogh is known as a Post-Impressionist, Andy Warhol for Pop Art, Banksy for Street Art, while Rodin and Brâncuşi are known as modernist sculptors. This categorisation process offers many benefits, such as making comparisons between artwork and artists more straightforward, for example, comparing Georgette Chen to Vincent Van Gogh is easier than to Banksy. Styles can also affect how a piece of art is aesthetically judged, placing rules onto the process, potentially helping people understand and describe their aesthetic preferences.

In art generation systems, the style in which the items are created is often not well defined due to the system itself being minimally constrained, searching through all possible permutations of colours on a canvas, meaning any style of artwork may be produced. A person's aesthetic preferences only represent a minute portion of the items which can

This article is part of the topical collection "Evolutionary Art and Music" guest edited by Aniko Ekart, Juan Romero and Tiago Martins.

✉ Edward Easton
eastonew@aston.ac.uk

Anikó Ekárt
a.ekart@aston.ac.uk

Ulysses Bernardet
u.bernardet@aston.ac.uk

1 Computer Science, Aston University, Birmingham B4 7ET, UK

potentially be created, the size of the search space compared to the region containing an individual's preferences makes finding aesthetically pleasing artwork difficult. Two methods exist for navigating to the region containing someone's preferences: using formalised measures to automatically navigate or asking the person themselves to direct the search. Formal measures are often limited in their effectiveness to accurately describe the aesthetic appeal of artwork and on the other hand, a human user, who despite being instinctively more proficient at assessing the aesthetic appeal, hinders the search in other ways, such as through user fatigue [41]. Another issue is a system's potential inability to reach certain areas of the search space, due to the choice of genotype representation, or design of genetic operators, effectively meaning the system may be unable to locate certain interesting items.

Attempts have been made to address user fatigue by giving the user more control over the direction of the search [16] or allowing the generation of artwork to be shared across multiple users [39]; however, these methods also have their limitations. The most effective method would be to constrain the search space to include only the items someone finds aesthetically pleasing, but due to the lack of exact formalisms for individual aesthetic preferences or even a full definition of which aspects may contribute to an aesthetic judgement, this would not be possible. Another way the search space can be constrained is to apply a specific style, this can make the search for more interesting items easier and consequently reduce user fatigue and improve the understanding of and ability to measure different aesthetic aspects. Examples of systems that have tried this approach include generating artwork in the Pop-Art style [8] or within the style of Mondrian [15].

This paper extends the versatile Axial Generation Process (AGP), first introduced in [17]. This system produces artwork within the Concretism style, either abstract images in two dimensions or digital sculptures in three dimensions. The process is capable of generating a wide range of visually varied artefacts of both types, the base process focused solely on the composition of the created items and did not involve colour or different shapes to make up the items. However, this approach limited how the AGP could be assessed, as formal measures relying on colour could not be applied. The assessment was instead restricted to using only a small number of form-based 3D measures. This left several unanswered questions, relating to the ability of the AGP to perform well across a wider range of measures as well as how the AGP compares to other generation systems at producing interesting items. Answering these questions required the AGP to be substantially extended to include calculating shape and background colour, allowing the content to be created using additional shapes, allowing these shapes to change size dynamically and allowing the content to be rendered in two dimensions from multiple viewpoints. Through implementing these changes the process is not only capable of generating a wider range of visually varied items but can also utilise a wider range of formalised measures to assess the output, allowing the AGP to be directly compared against other generation systems. The remainder of the paper is organised as follows: "Existing Generation Methods" provides a brief overview of existing 2D and 3D generation methods, "Choosing a Style" looks into how a style was chosen for the system and the benefits of doing so, "Axial Generation Process" details how the Axial Generation Process generates its content, including the extensions to the process, "Evaluation of the Axial Generation Process" provides an evaluation on how effective the process is at generating interesting content based on a variety of measures and these results are then compared to two existing systems in both 2D and 3D: a particle-based generation system [10, 16] and a pixel/voxel-based generation system [11, 20, 22, 24], finally the conclusions are discussed in "Discussion and Conclusion".

Existing Generation Methods

A wide variety of methods have been used for generating artwork within Evolutionary Art systems [14, 20, 28, 39, 44, 46], these different processes can be split across multiple categories. For the purpose of this work, the categorisation into methods that generate 2D and 3D items is the most relevant.

2D Items

Two-dimensional items are widely studied and the generation of images based on mathematical expressions is the most popular approach. These can be represented by lisp-style expressions [14, 40] or in the form of expression trees [10, 20, 24, 25, 31, 46]. The way these expressions are used to generate an image takes a variety of forms, the most common procedure is to use a value generated by the expression to set a property of each pixel in an image, such as the luminosity [21, 25], using multiple expressions to calculate the RGB value of a pixel [11, 20, 24] or mapping the value to a colour lookup table [14]. Expressions have also been used to calculate the position and colour of a line over a series of time-steps [10, 16] and to create animated content [31, 44]. Other underlying data types are routinely used, such as Compositional Pattern Producing Networks (CPPNs) [23, 33, 39, 44]; Context-Free Design Grammars (CFDGs) [46]; encoding the parameters to use in another generation process [12], using SVG [15], or maintaining a list of shapes to be placed on the canvas [9]. Entirely custom representations have also been designed to solve particular problems, for example, to efficiently hold the data to create a piece of art in the style of Mondrian [8]. In a similar sense the AGP generates within a specific style; however, unlike some of the existing systems which attempt stylistic generation, the main focus of our system is not the style itself but how a style can be applied to improve the output. Custom representations are also avoided, with the more common representation of expression trees being used. In addition to this the system primarily works in 3D using this as a basis for generating the 2D content.

3D Items

3D item generation does not have the same number of existing systems; however, there is still a diverse range of representations and methods being used. This includes manipulating the colour and visibility of voxels using CPPNs [22] or Context-Free Grammars [2, 29, 36], using Graph Grammars to generate a set of points in 3D space [30], using Shape Grammars to create 3D items [5, 32, 35] and evolving parameters to use within an external generation system [28, 34]. Other non-direct generation methods have also been used, such as by taking 2D content and adding the extra dimension [18]. The process presented in this paper works in the opposite direction, converting from 3D to 2D allowing the same content to be generated without losing any information about the original item.

Choosing a Style

Often the style in which content is generated is an after thought for generation systems. As a consequence, the style of artwork produced in Evo-art systems is often underdefined, being placed under the wide ranging banner of abstract art. Categorisation of the generated output provides numerous advantages when assessing and understanding the items being generated. First and foremost the style of a piece of art influences how it is judged aesthetically, for example, Post-Modern sculpture will use a different set of criteria than classical sculpture. Without an understanding of the style which is being created, it becomes more difficult to assess the content, even by people, who can use a style as a reference point for whether a piece is aesthetically pleasing or not. Abstract art has numerous sub-genres contained within it. During the process of creating the system [17], one of the sub-styles nicely described the content which was being created: Concretism. As defined by Tate: “Concrete art is abstract art that is entirely free of any basis in observed reality and that has no symbolic meaning” [42], Concretism was introduced in 1930 by Theo van Doesburg who released the *Manifesto of Concrete Art* [45], outlining the vision for Concrete art, through six well-defined tenets describing how and why Concrete art should be created. Due to the style being disinterested in describing natural objects or sentiments, it fits nicely with the auto-generated nature of Evo-art systems. The similarities continued to grow through further investigation of the manifesto. The third tenet describes that concrete art ”must be entirely built up with purely plastic elements, namely, surfaces and colours” [45], this has a striking resemblance to the use of basic shapes to generate artwork allowing a connection between Concretism and the AGP to be established. The style was then invaluable not just in how the generated content could be assessed but it was also used to influence the system and how it could be improved and grow.

The remaining tenets also fit well with the constraints that the AGP applies to its generation process, albeit with a slightly amended understanding of terminology to include computational-based generation. The second tenet describes two aspects, the first is how the work of art should be conceived entirely in the mind of the artist before being created, this is analogous to how an Evo-art system maintains all images within the genotype before rendering them into the phenotype. The second aspect of tenet two describes how the artwork should not be influenced by external sources, for example, nature, sentimentality or emotional content. This once again fits well with the AGP, and Evo-art systems in general, as no external influences are used to create the artwork, any references to nature, emotions or other concepts are most likely accidental.

The fourth tenet describes how the artwork should be constructed, through using simple processes. To apply this, it needs to be considered outside of its likely original intention. Within the context of computer generated artwork, highly complex items are created from basic shapes, such as squares, representing the limitations imposed by graphics cards and screens. In this sense, simple construction includes primitive shapes, such as spheres, circles, cubes, squares and triangles, within this understanding of the tenet, it is applicable to Evo-art systems.

The fifth tenet considers how the artwork should be generated, looking for exact techniques. This tenet is easier to apply from a computational perspective as by their nature, computers already use mechanical techniques for all tasks. Finally, the sixth tenet asks the artwork to aim for clarity, this is not always followed by the AGP, due to other considerations about the system including improving the variety of the content the system can generate, this requires the system to generate a wide range of distinct content, which includes the more simple artwork the sixth tenet is aiming for as well as the more complex. However, by amending generation parameters a more strict generation of Concrete artwork can be achieved, see Fig. 1.

The understanding and application of these tenets within the Evo-art context helps to achieve the aims for this system. Applying these constraints leads to a smaller search space to

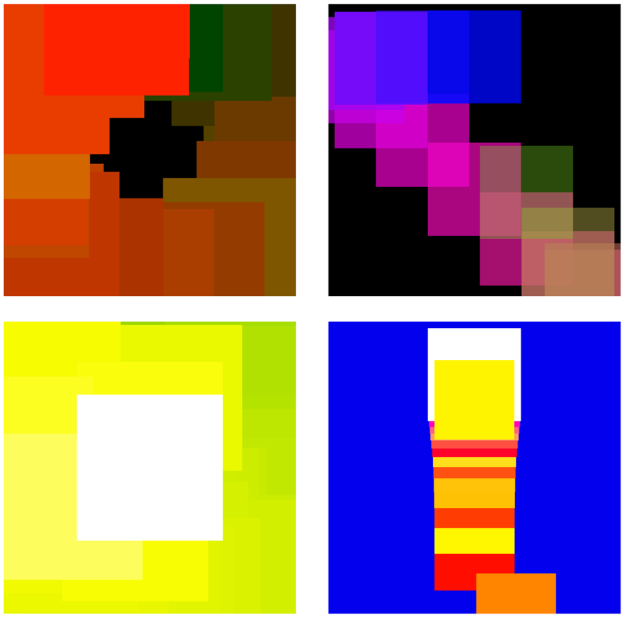

Fig. 1 Generated images which adhere more strictly to the tenets of Concretism

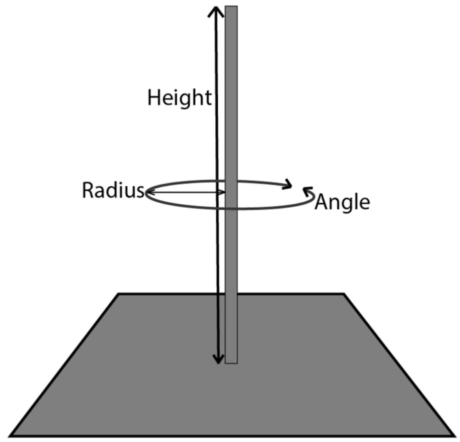

Fig. 2 The use of the height, the radius and the angle to place the geometric shapes around the axis

Algorithm 1 Base process

```
totalItems ← n
heightTree ← Tree
angleTree ← Tree
radiusTree ← Tree
Points[] ← [totalItems]
for k ← 1 to totalItems do
    y ← heightTree.getValue(k)
    angle ← angleTree.getValue(k)
    radius ← radiusTree.getValue(k)
    x ← radius × Sin(angle)
    z ← radius × Cos(angle)
    Points[k] ← (x, y, z)
end for
```

be traversed allowing more interesting artwork to be located more easily. It allows a deeper understanding of the content being generated and how it may be assessed, potentially making describing the aesthetic features an easier process.

Axial Generation Process

The base Axial Generation Process (AGP) creates items by combining a specified number of geometric shapes within set bounds, three values are used to position the shapes around a central axis: the height, the angle and the radius. Figure 2 shows how these three values are used to represent a precise location in 3D space. To calculate these three positions for each shape, the base process maintains three expression trees. The expression trees used differ slightly from other implementations, which commonly use position details and instead only use two possible values for the leaf nodes, the index of the shape being placed and a random value between 0 and 1, as the position data for each item is not known until after the expression has been evaluated. The base process, shown in Algorithm 1, is used to position the shapes for both the 2D and 3D items.

Eight parameters are used to control this base process, described in Table 1. These include the number of shapes to be added, the bounds the items should be placed within and how the items should be rendered. Amending these parameters can change the visual structure of the generated items as shown in Fig. 3, which contains items generated using only the base process.

The base process [17] serves as a proof-of-concept, showing that the method can generate a wide range of items, both in 2D and 3D. The base process provides many areas which can be extended to further increase the range and diversity of items that can be generated allowing a wider range of measures to be used to assess the created content while keeping within the bounds of Concretism. We have extended the base method with four new features detailed below: (1) colour, (2) dynamic sizing, (3) different base shapes, and (4) multiple viewpoints.

Adding Colour

Many factors contribute to making an aesthetic judgement, One very common factor is the colours contained within the artwork. Colour is a major component within art and is

Table 1 Parameters used by the AGP

Total items	The number of items which will be placed in the artwork
Maximum height	The furthest distance away from the base the items can be placed
Maximum radius	The furthest distance from the axis the items can be placed
Is 3D	Indicates whether 2D or 3D items should be generated
Element size	The radius of the circle to be placed on the canvas (2D only)
Background colour	The canvas colour to be used when drawing the image (2D only)
Canvas width	The width of the bounds to render the items in (2D only)
Canvas height	The height of the bounds to render the items in (2D only)

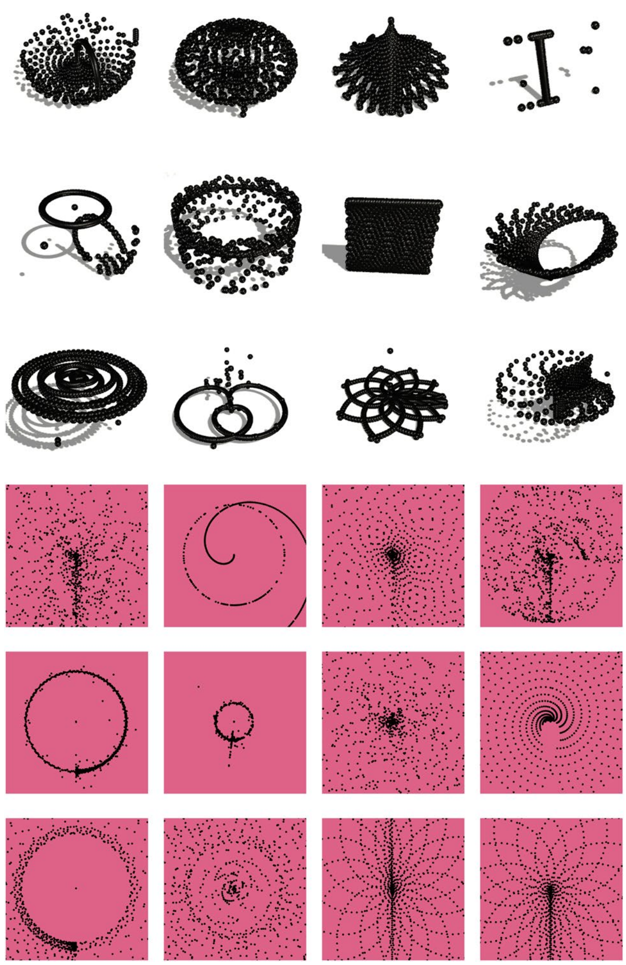

Fig. 3 A variety of generated items using the base AGP

the focus for a large number of existing generation systems. This is also important within Concretism, being specifically mentioned in tenet three. Due to its importance to aesthetic judgement, it also forms the basis of a large number of existing measures which attempt to assess images [1, 27, 38], looking at aspects such as the relative positioning of the colours or which gradients are present. Due to the stalwart nature of colour as a part of aesthetic judgement, adding colour into the generated sculptures and images was a necessary addition. The colour calculations rely on three expression trees representing the R, G and B components of each shape and additional parameters were introduced to allow greater control over how the AGP handles the colour, as shown in Table 2. Examples of axial artwork created in full colour are shown in Fig. 4.

The background colour is another important feature and can also be a significant factor used when assessing artwork, with the brain amending its response to seeing a colour based on the background it is presented against [48]. It can be used to emphasise or hide different parts of the form of the artwork, dictating the importance of including the background colour in the AGP. The background colour is only implemented for 2D items as it is inappropriate to implement in 3D. The environment a 3D item is viewed in is context-specific and based on how the sculptures may be used, it also includes complex considerations, such as lighting position making it out of scope. The background colour feature is demonstrated in Fig. 5, where the same forms are shown against different backgrounds, the images in Fig. 5a and e emphasise the form of the generated shape and the remaining images mask it, for example, through having the background colour match the colour of some of the shapes being placed, as seen in Fig. 5f–h, or through using contrasting colours in the background. Figure 5d has a yellow background making the items on the bottom left of the form not as noticeable as in Fig. 5a. Similar to shape colour, the background colour uses three expression trees which represent the R, G and B values and relies on one additional parameter in Table 3 and the background colour parameter from the base process.

Dynamic Sizing

The base process of the AGP focuses on the form of the generated items, part of this form depends on the primitive shapes which are being placed within the output. Within the base process, these shapes were limited to a single size revealing another extension that can be implemented to further increase the range of forms the system can generate. This aspect was amended in two ways, the first is to introduce the dynamic sizing of the individual shapes within a

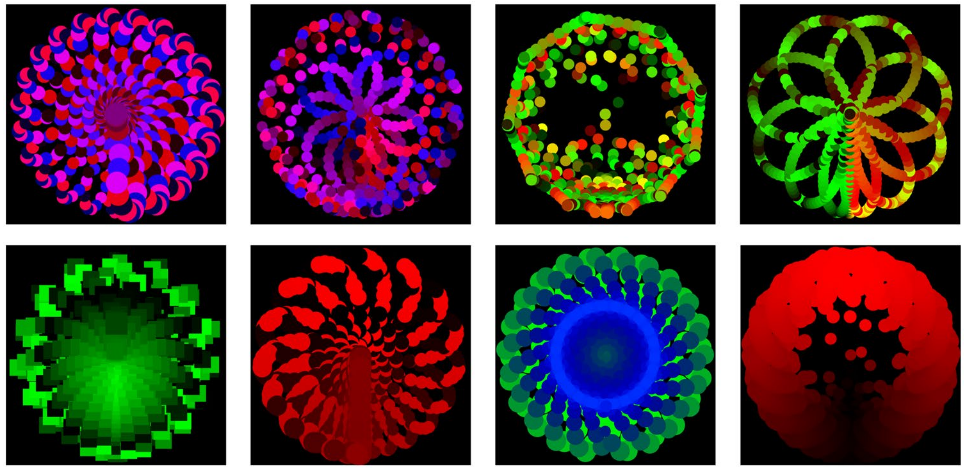

Fig. 4 A variety of colour images generated using the AGP

Table 2 Parameters used by the AGP relating to calculate colour

Calculate item colour	Whether the colour of the shapes should be calculated
Item colour	The default colour used for the geometric shapes
Standardise item colour	Whether all shapes should have the same colour

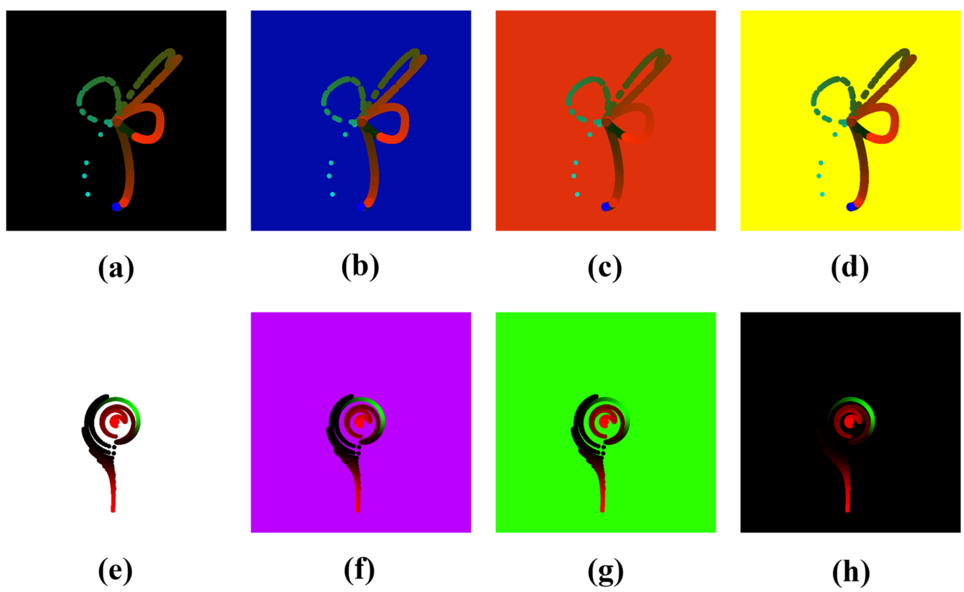

Fig. 5 2D images with various backgrounds generated using the AGP

Table 3 Parameters used by the AGP relating to calculating colour

Calculate background colour	Whether the background colour should be calculated

specified maximum and minimum value. This uses three further expression trees which calculate a value for the scale of the shape along the X, Y and Z axes, if the scale is standardised through the use of the parameters, the scale value for the x-axis is applied across all axes. Forms and composition of the items are the basis of Concretism, extending the range of forms available also benefits the coherence of the system to its style. The parameters which control the dynamic sizing are shown in Table 4. A selection of both 2D and 3D items, showing the capability of dynamic sizing, are shown in Fig. 6 which uses standardised scaling and Fig. 7 which shows shapes to be scaled independently

Table 4 Parameters used by the AGP relating to calculating colour

Maximum element size	The maximum width or height a shape can have
Minimum element size	The minimum width or height a shape can have
Standardise scale	Whether a shape should have the same scale across all axes

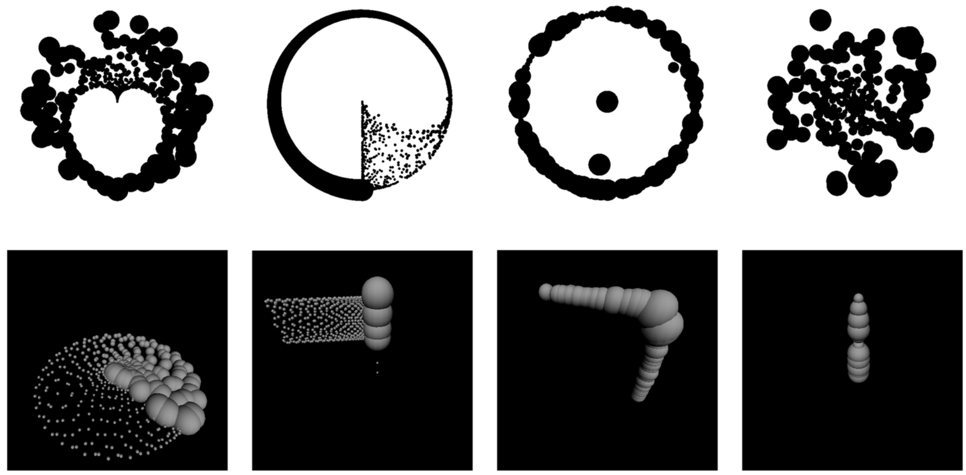

Fig. 6 Generated content using standardised dynamic sizing of shapes

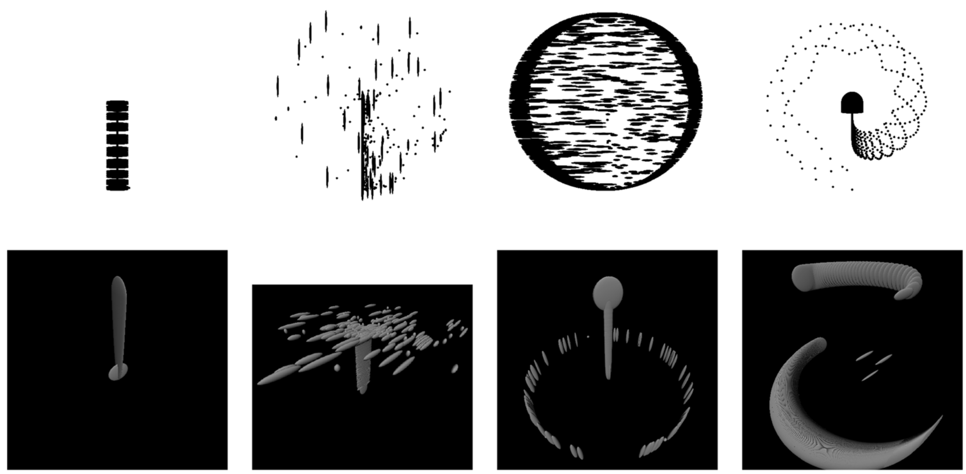

Fig. 7 Generated content using non-standardised dynamic sizing of shapes

on all axes, indicating the influence this change can have on the final output.

Shapes: Spheres and Cubes

In addition to allowing the shapes to scale, the process has also been extended to increase the range of primitive shapes used to build the items, allowing either a sphere or a cube. These two elements were chosen as they are ubiquitous across all graphics software, keeping inline with the discussion about tenet four, which suggests more complex base shapes should not be included. To keep this feature consistent with the process and allow the shapes used to be dependent on the evolutionary processes, a single expression tree is used to determine which type of element is used to create the item. To keep the performance of the process stable, only a single shape per output item is allowed; however, this greatly affects the look of the created items, as shown in Fig. 8. As with all the form-related extensions, the automatic specification of the shape can be turned off and a single element can be specified, using the two parameters in Table 5.

Rendering the Items from Multiple Viewpoints

The AGP abstracts the rendering process away from the generation method, allowing the same output to be rendered in either 2D or 3D. Rendering the 3D content is a simple process and involves placing the primitive shape at the specified location in 3D space. For 2D content, the rendering process works in a similar way to creating 2D art, where a 3D object

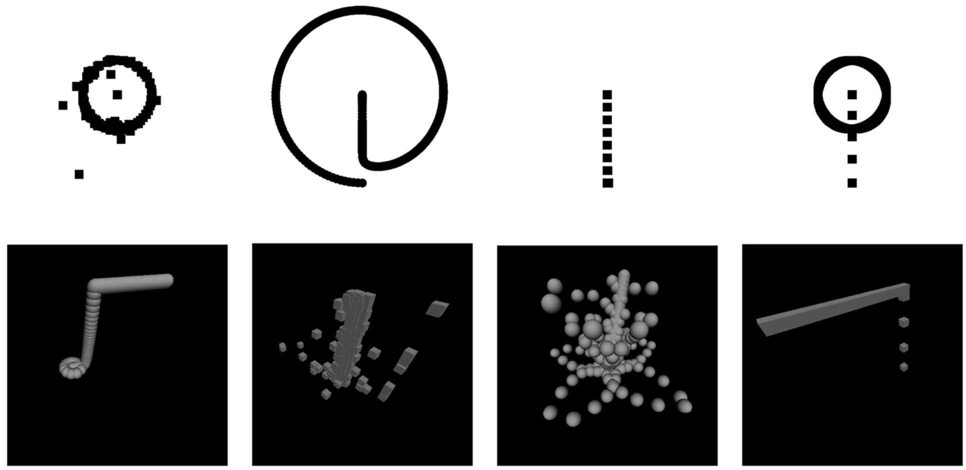

Fig. 8 Generated content using different types of base element

Table 5 Parameters used by the AGP relating to set the base element type

Generate shape	Whether the shape should be automatically set
Base geometry	The shape to use if not calculated automatically

is abstracted down onto a 2D plane. The base process only allowed this flattening to happen from the top–down view of the items, with the size of the items determined by their distance from the canvas, referred to as perspective-driven flattening [17].

As both 2D and 3D content initially gets created as a 3D point cloud, this opens up a further extension for the process: allowing the 2D items to be rendered from multiple viewpoints. This feature is incompatible with the perspective driven flattening as the distances to the canvas can become very large causing single shapes to cover the entire canvas and so the perspective calculations have been removed.

Six additional viewpoints can be used to render a 2D item from the directions shown in Fig. 9. The viewpoints are limited to six views to find a balance between increasing the range of items that can be created and keeping the difficulty of using the process as low as possible. Allowing different viewpoints affords two major benefits which will be demonstrated further in the "Evaluation of the Axial Generation Process", where it increases the range of 2D items the system can create as well as allowing 3D content to be assessed using 2D-based measures, where a snapshot is taken from each viewpoint and assessed using the 2D measure. The viewpoint is controlled by a single parameter allowing the viewpoint to be specified from the choices of Top->Bottom, Bottom->Top, Left->Right, Right->Left, Front->Back, Back->Front, Fig. 10 shows the same item from different viewpoints, indicating the importance of the viewing position on the judgement of the sculpture, giving additional focus to the form of the item. The chosen

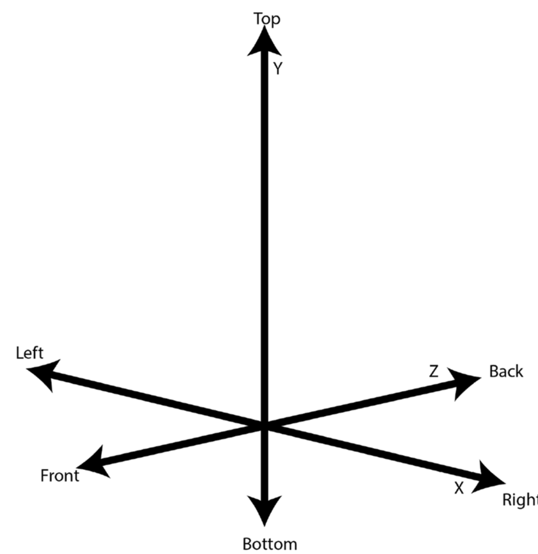

Fig. 9 The directions for each viewpoint

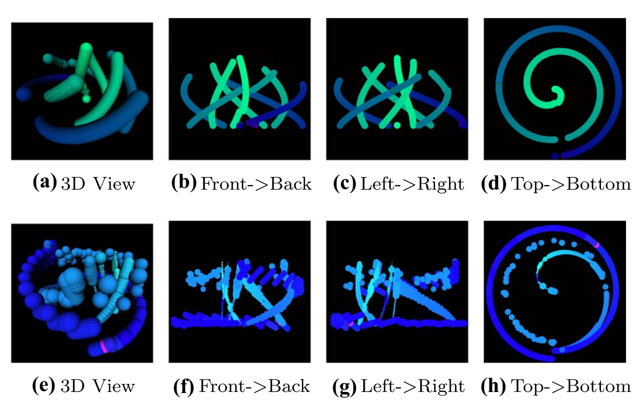

Fig. 10 Different views generated for the same 2D item: **a** and **e** display the item in 3D and **b**–**d** and **f**–**h** the resulting views, respectively

viewpoint can also dictate the overall shape of the form, where viewpoints that use the Y-axis as the height create square forms rather than circular ones.

Figures 11 and 12 display a range of items created using either the base process (top half) or taking advantage of all the extensions (bottom half) in both 2D and 3D. As can be seen in the items which use the extended features, the forms are much less regular, this is related to the dynamic sizing of the shapes which can allow different patterns to be expressed. The element colour also adds a significant amount of diversity, where even sculptures that use similar shape sizes can appear dramatically different. This highlights the effectiveness of these five extensions to the AGP representing a significant improvement allowing the process to generate a wider range of more interesting items.

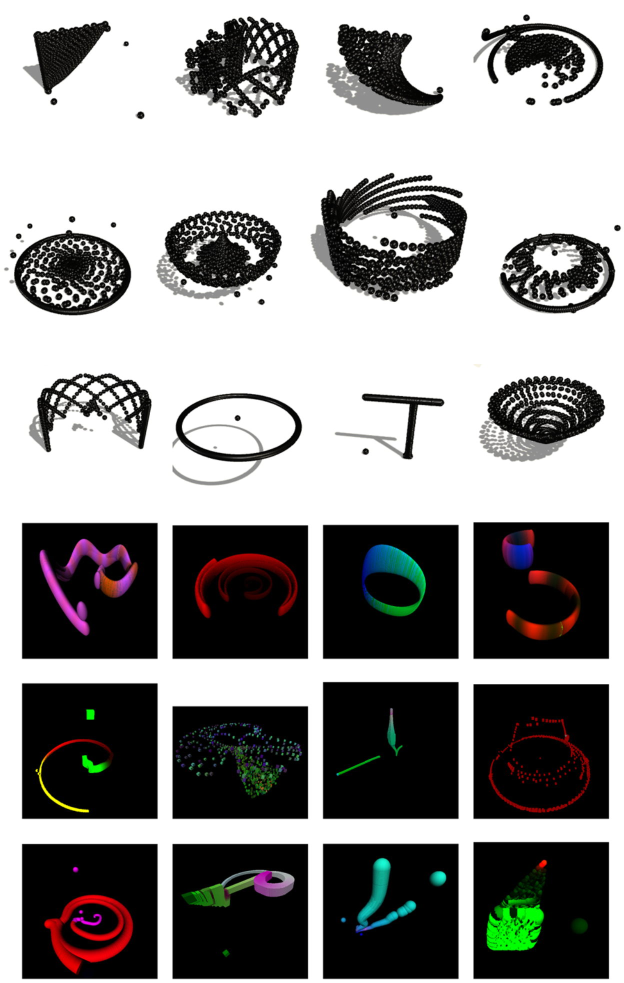

Fig. 11 A variety of randomly generated sculptures using the AGP

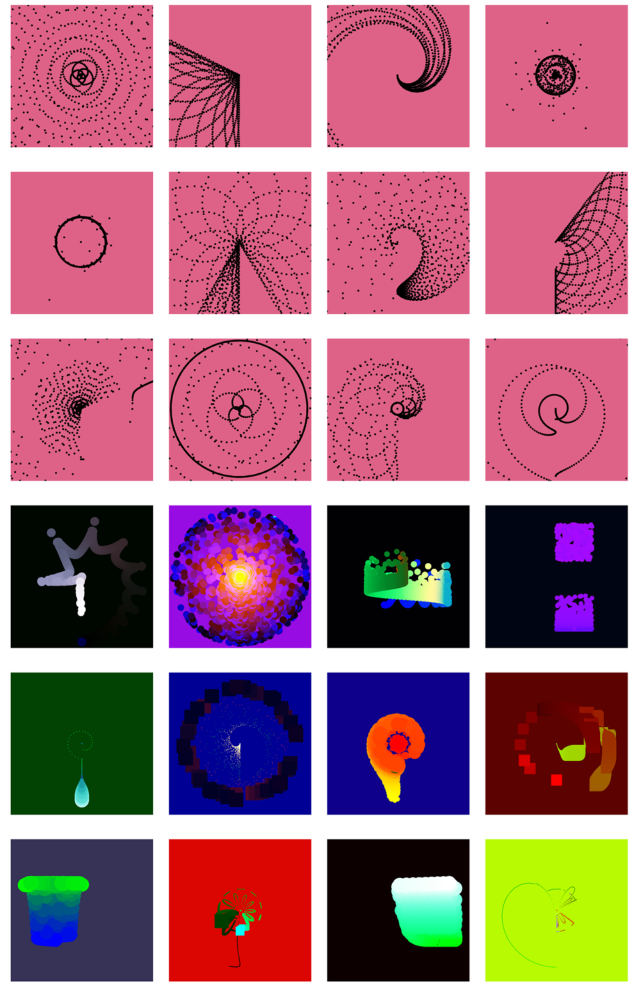

Fig. 12 A variety of images generated using the AGP

Evaluation of the Axial Generation Process

In order for the AGP to be useful within an Evolutionary Art context, it needs to produce artefacts of interest. A common method is to classify artefacts based on their "novelty" and "value" [4, 6, 47]. However, novelty is a context-specific measure and cannot be directly assessed without knowing how the output might be interpreted or used. Alternatively, the process can be evaluated for its versatility, i.e., the range of items that can be generated which show a variety of values across different measures. There are many formal measures that aim to describe the aesthetic appeal and can be used to evaluate these items. These measures are ubiquitous within automated and semi-automated Evo-art systems; however, they do not represent a full aesthetic judgement of an item. Instead they only represent some of the properties of the output, representing contributing factors for aesthetic judgement. These measures only provide a limited insight into the aesthetic appeal of the generated artwork; however, they form a solid base upon which further evaluation can be completed. This includes empirical evaluation which allows a more complete assessment of the aesthetic appeal of the generated items. Due to the limitations of the base implementation of the AGP not including colour, the set of applicable measures to assess the output was limited. The chosen measures also had to be applicable to both 2D and 3D items. Due to these restrictions, established measures, such as the Global Contrast Factor [27], the colour gradients present within an image [38] or the naturalness of the image [1] were not considered. Instead, three canonical properties are used, all of which are considered to contribute to the aesthetic appeal of a generated item: the level of symmetry, complexity and compressibility. These measures were chosen for their relevance to aesthetic judgement, rather than their specific presence within Concrete artwork. This allows the initial viability of the AGP to be more thoroughly assessed on whether it can generate a wide range of items. To our knowledge no measures exist which consider the aesthetic appeal of Concrete art, the use of which would potentially not be applicable to other systems making the evaluation less meaningful.

First, symmetry is known to be an important factor in the aesthetic judgement [43], the AGP places items around a central axis, often yielding highly symmetrical items. From the many available methods to calculate symmetry, the measure presented in [19] is used as it calculates the level of symmetry for point clouds and, therefore, is applicable for both 2D and 3D items. This measure works by finding multiple candidate symmetry planes and then refining their positions using the Levenberg Marquardt algorithm that minimises the error between points reflected in the plane. If the error is lower than a specified threshold, the item is considered symmetrical around it. The score calculated for each item is the average distance error across all detected planes for an item.

Second, the complexity of an item is also known to impact the aesthetic appeal [3, 26]. The method employed here is based on the measure utilised by Birkhoff [3], involving counting the vertices in an item. Instead of fully rendering each item and counting the vertices, the value was estimated by merging items with similar positions, using a distance threshold of 0.2. The remaining geometric items in the artwork are then counted and divided by the original total, the intention is if items have many shapes in the same position, they will be less complex.

Finally, we use the Global Normalised Kolmogorov Complexity [37], shown in expression (1), calculating the compressibility by encoding all points into text and then compressing this string using ZIP compression, similar to the process used in [20].

To perform the assessment, 500 items in both 2D and 3D were generated using the AGP, without any extensions applied a selection of generated items and their calculated values are shown in Fig. 13:

$$\frac{\text{originalSize} - \text{compressedSize}}{\text{originalSize}}. \quad (1)$$

As shown in Fig. 14, both 2D and 3D items were generated within a wide spread across the ranges for both symmetry and complexity; however, the process was unable to generate items that simultaneously had a high level of complexity and symmetry. In addition, there were no compressibility values lower than 0.6067 for 3D items and 0.6747 for 2D items, indicating that the lower end of the spectrum could not be reached. These two gaps potentially represent a limitation with the base process when attempting to generate novel items, as these regions of the search space may not contain a large number of artefacts.

Filling the Gaps

To determine whether items exist within the gaps in the search space (i.e., no items with both high complexity and symmetry and no items with low compressibility), four evolutionary algorithms were used. These attempted to maximise both symmetry and complexity as two objectives and minimise the compressibility as a single objective, for both 2D and 3D items. The evolutionary parameters are outlined in Table 6. A visual inspection of the resulting distributions, displayed in Fig. 15, shows that the evolutionary optimisation was not able to fill the gaps in the measured values. The compressibility measure did not reduce any further than the original values, although more items fell within the lowest band and the symmetry and complexity could not be simultaneously improved. There are two possible explanations for this: the process itself is potentially limited and unable to produce both complex and symmetrical items or the measures themselves do not accurately represent the details of the generated artwork, highlighting the potential need for more appropriate formalised measures to be created.

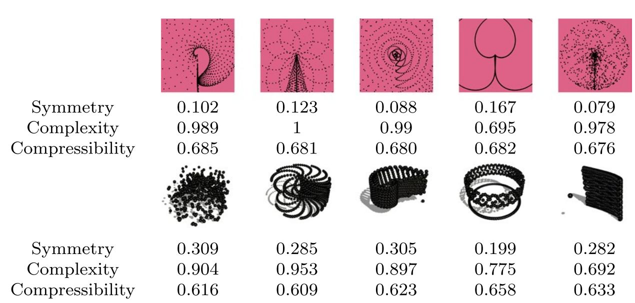

Fig. 13 Example output and associated values

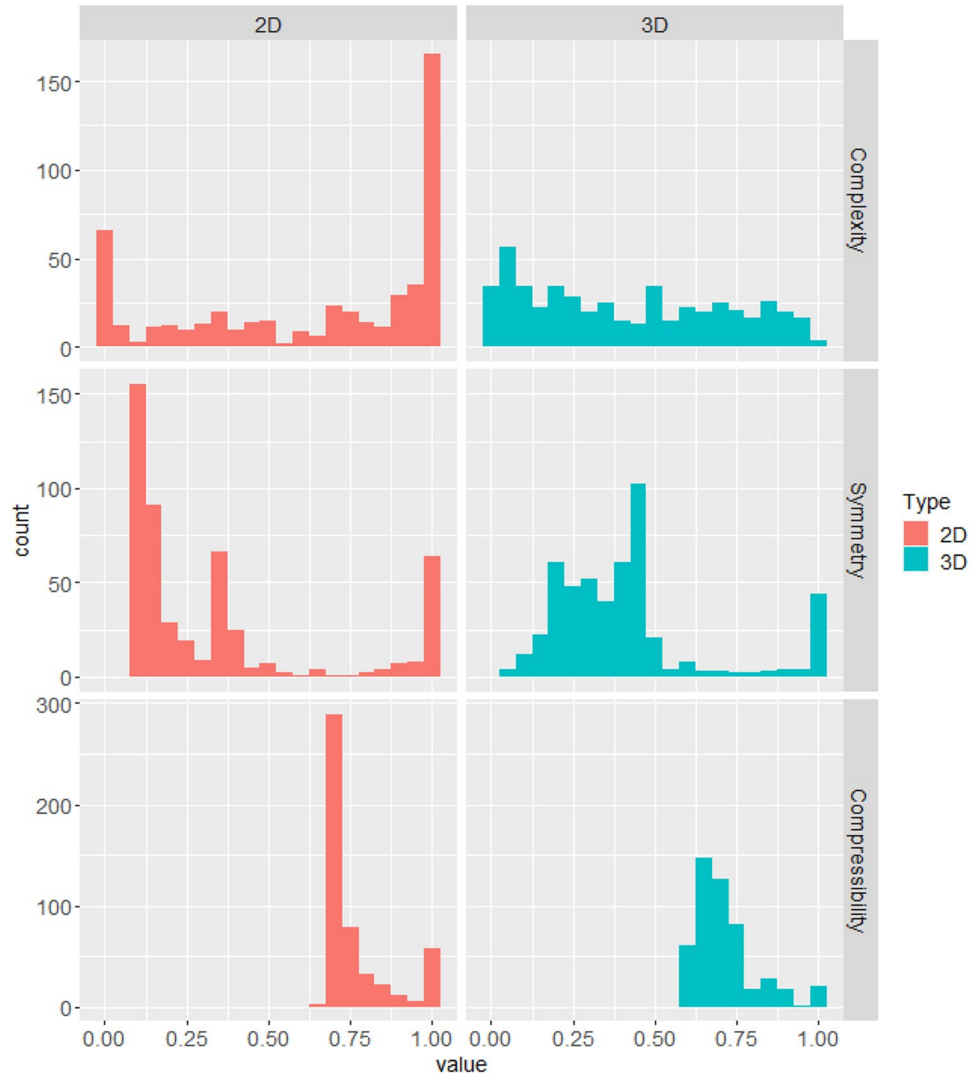

Fig. 14 The distribution of the randomly generated items

Table 6 Evolutionary Algorithm parameters

Total generations	30
Mutation probability	0.7
Initialisation method	Full
Selection method	Tournament (k=3)
Population size	150

To ensure that the process itself was not the limiting factor, the extensions were applied and further analysis was performed. To determine whether the extensions helped to generate interesting and novel items, the generated content was compared to two well-established systems.

To compare both 2D and 3D, two 2D processes were chosen that while not originally designed for 3D, only required small changes to allow 3D items to be generated using their distinct features. The first additional method introduced is a pixel-based method that calculates the colour of each pixel within an image for 2D items. To amend this to work with 3D items, an additional depth parameter was added, allowing 3D items to be created, similar to the implementation in [22], where voxels are used to build sculptures with variations in form achieved by calculating the opacity of each voxel. Examples of items generated using this process in both 2D and 3D are shown in Fig. 16.

The second method is an updated version of the process used in [10], this process tracks the position of a specified number of particles across a canvas over a specified number of time steps, drawing a coloured line between the start and end position for each time step and then blurring the entire image. This process has been used within 3D before [18], however, that implementation focused on allowing a user to explore artwork in Virtual Reality (VR) with the 3D aspects were constant values based on the particle index. For this analysis, the depth is implemented via a new parameter added onto the genotype which then allows the paths of each particle to be traced in three dimensions. Due to

Fig. 15 Distribution of the items after the evolutionary run

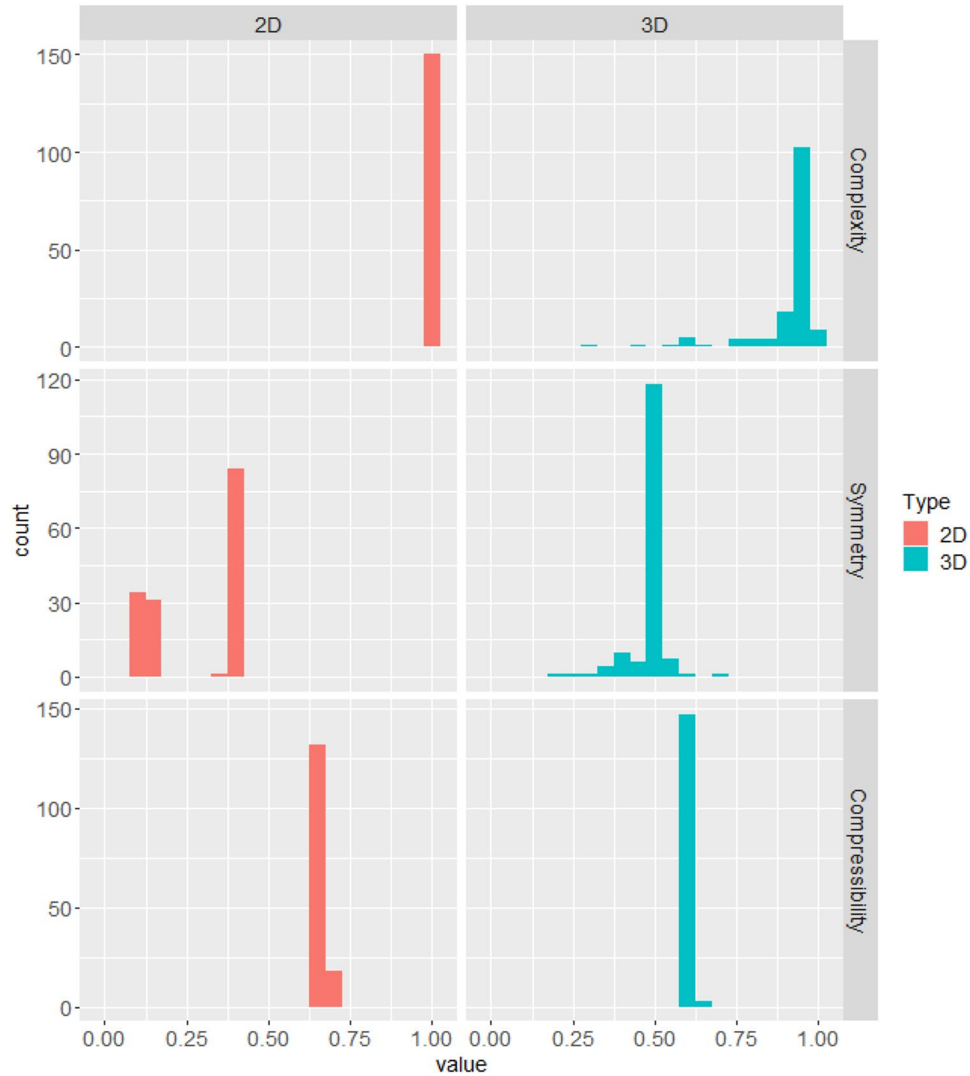

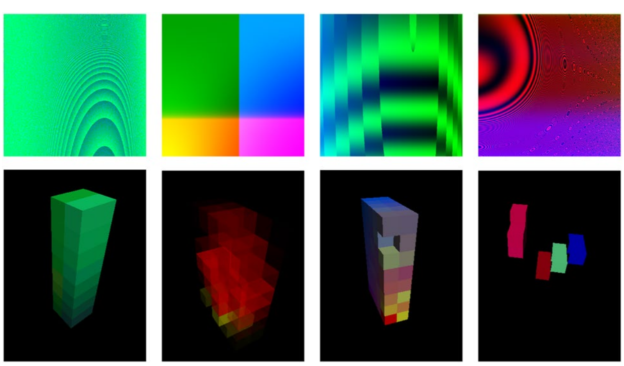

Fig. 16 Pixel/voxel-based generated content

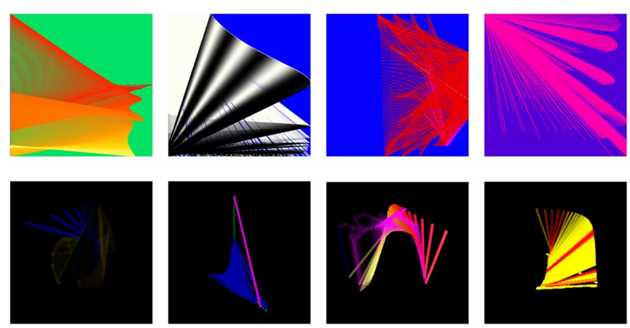

Fig. 17 Particle-based generated content

the extension of the process from 2D to 3D the blurring which originally happened after each time step was no longer included as it would be overly complex for 3D items to implement in a similar manner to 2D items. Examples of generated content using the particle-based process are shown in Fig. 17.

To further assess the items generated by the AGP and the two contrasting processes, different measures are implemented. These follow the same themes as for the base process, considering the symmetry, complexity and compressibility, but choosing measures which are sensitive to both colour and form, ensuring all the applied extensions to the AGP are taken into account and allowing a fair comparison between all three processes. For 2D items, the final generated image was used for the evaluation taking into account the varying shapes and colours which are present in items generated through all three processes. To assess the complexity of the 2D items, Shannon's Entropy [37] was used, which calculates how much colour information is present in the image. The symmetry was calculated using the SYM4 method [13], which compares the luminance values of each pixel across 4 different quadrants of an image and finally, the Global Normalised Kolmogorov Complexity [37] was calculated by first encoding the item into a bitmap and then applying ZIP compression.

The implementation of these measures differs slightly from the original, which consider every pixel in the image as having the potential to contain information. For this assessment, the background of the images was set as transparent and only pixels which had a calculated colour value were used to normalise the values obtained. This decision was made for three reasons: first the pixel-based systems do not have the concept of a background making the comparison less fair between the three generation methods, second, measures such as SYM4 can have their values artificially increased if large blocks of a single colour are present and finally, it allows an increased focus on the form of the item being assessed. 500 2D items from each process were randomly generated, using a max tree depth of 5 with the distribution of the values shown in Fig. 18 (please note the logarithmic scale) and examples of the generated content shown in Fig. 19, indicating the differences in blank space between each generation method and how the compressibility generally occurs on a limited range.

All three generation processes created content exhibiting a wide range of values for both the complexity and symmetry measures, with all processes showing a relatively consistent spread for both measures. The AGP produced content that generally had higher values for the symmetry, this is expected due to the rendering process of the AGP placing

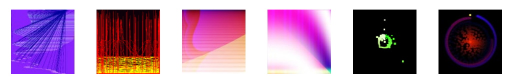

Symmetry	0.59	0.82	0.36	0.64	0.23	0.99
Complexity	1.66	0	4.69	3.56	3.96	4.14
Compressibility	0.97	0.96	0.81	0.73	0.99	0.98

Fig. 19 Example 2D output across all processes

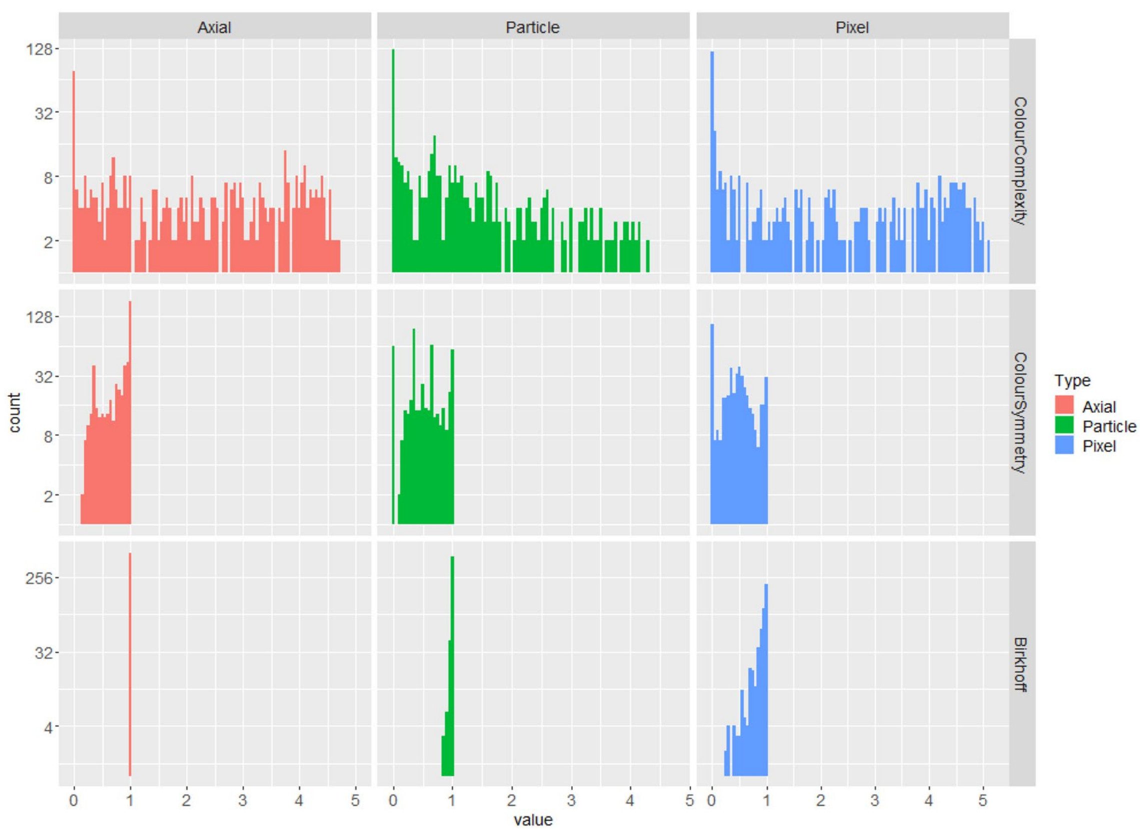

Fig. 18 The distribution of the 2D randomly generated items across all generation types

items around a central axis. In addition to this, the AGP also has a slightly more uniform spread of values for the complexity measure. The compressibility is the only measure, where significant differences exist between the processes, where the pixel-based method covers a much wider range of values than either the Particle or Axial processes. This is potentially due to differences in the generation process, where in both the Particle and Axial processes, the focus is on generating the form of the item within the bounds of a canvas, which often leads to larger amounts of blank space being present in these images than there is within the Pixel process.

For the 3D comparison, the symmetry and complexity for both the colour and form of the items were calculated complemented by the compressibility as calculated in the base process evaluation. To assess the form of the generated items, the symmetry method proposed in [19] was used, and the form complexity was calculated using the viewpoint entropy [7]. The viewpoint entropy is used to determine the best viewpoint of a 3D scene by calculating how much of the scene is visible from each viewpoint. For this evaluation, the six viewpoints available in the system are used to calculate the value and the average is provided across all six viewpoints, indicating how much of the sculpture is visible.

For the colour-based measures, the symmetry was calculated using a similar method to the form symmetry except the final score depended on how close the intensity of the colour was for each matched point across the plane of symmetry. An additional constraint was also added when detecting the original planes of symmetry: if matched item's colours were more than a specified threshold away from each other, this plane was discarded. The average colour error across each detected plane of symmetry is the colour symmetry value presented. The colour complexity follows the same process as the viewpoint entropy except Shannon's entropy is calculated for each viewpoint and the average value provided. As with the 2D items and due to the background colour not being calculated for 3D items, the background was once again ignored for calculating Shannon's Entropy and only pixels containing a set colour were included in the calculations. In addition to this, due to the scales of the 2D and 3D items not being the same (2D distances are measured in pixels and 3D are measured in metres) the 3D items needed to be scaled to ensure they entirely covered the 2D canvas without losing any information. 500 3D items for each of the processes were generated, using a max tree depth of 5, providing the results as shown in Fig. 20 with examples of the 3D content shown in Fig. 21.

The AGP showed the greatest variations across the measures, followed by the particle generation process. For all five measures, the Axial process generally had the widest and most uniform spread of values. For the symmetry measures, this is expected due to the shapes being placed around a central axis, however, for the colour and form complexity this result is more surprising, indicating that the AGP is capable of creating items that show a wide range of information. The voxel-based system produced the lowest range of values across most measures; however, this can be explained by how the items were generated, where they were all cuboids composed of voxels, this means that the symmetry would be identical across all the items and the number of objects visible from each view would also be constant, explaining the low range of values this process generates. One interesting aspect all processes demonstrated is shown in the colour complexity, where the highest number of items had the lowest value of between 0 and 0.2. This is due to the initial random generation creating a high number of uninteresting and uncomplex items. This is an aspect, where the AGP performed slightly better than the Voxel and Particle processes by creating a lower number of these types of items.

Finally, the compressibility encountered the same issues as with the base process evaluation and 2D evaluation, where only a small range of values are created across all three generation systems. Once again, due to the increased variability in how the Axial and Particle items are stored, it would lead them to have a higher range of values than the voxel-based system whose definitions could only ever differentiate based on colour, leading to higher levels of redundancy within the files, leading to consistently high values being found.

Conclusions and Discussion

The extensions to the AGP presented in this paper all have a single focus, to improve the versatility of the Axial Generation Process, based on the results presented in [17]. The process itself applies the Concretism style to constrain the search space to improve its ability to generate interesting and varied content. The base process was capable of generating content which exhibited a wide range of values across symmetry and complexity measures; however, issues were identified and it was not clear whether this meant the AGP was limited or how these results compared to existing systems. To answer these questions and increase the range of content generated by the AGP, several extensions were applied to the process and a more extensive evaluation performed. In general, the AGP produced items in both 2D and 3D which displayed a wide distribution of values across measures representing the symmetry, complexity and compressibility, indicating that it is capable of generating a wide range of visually varied, complex and interesting images as well as digital sculptures, comparing favourably to two well-established systems.

The identification and application of a specific style was beneficial as it provides a benchmark through which

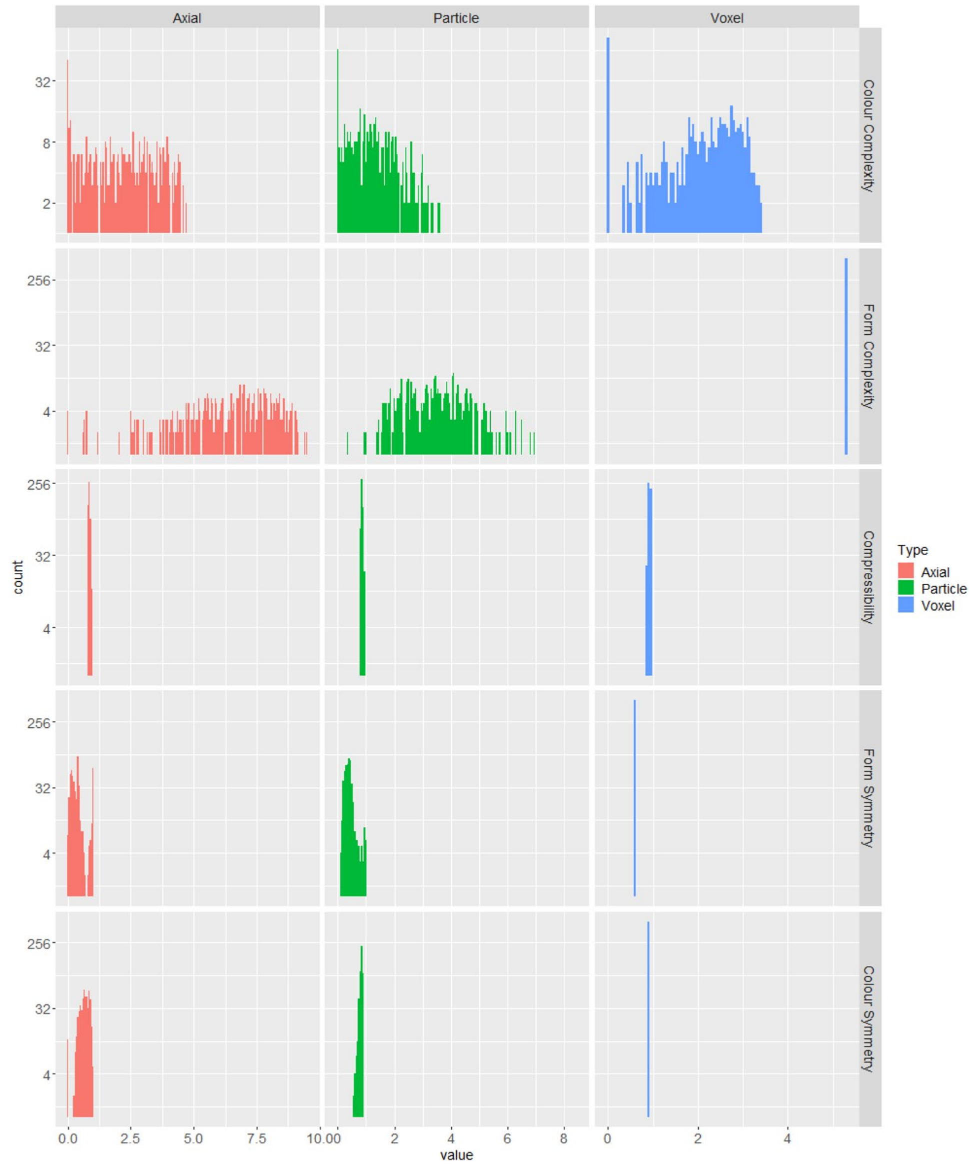

Fig. 20 The distribution of the 3D randomly generated items across all generation types

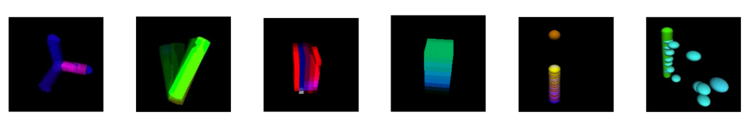

Fig. 21 Example 3D output across all processes

Form Symmetry	0.88	0.29	0.62	0.62	0.42	0.42
Colour Symmetry	0.81	0.79	0.92	0.92	0.74	0.60
Form Complexity	1.21	4.71	5.31	5.31	5.08	4.58
Colour Complexity	0.37	1.75	2.17	1.98	3.38	1.63
Compressibility	0.86	0.83	0.91	0.91	0.89	0.66

the content can be judged and any aspects, which help to form this judgement, can be better understood in terms of the style. Concretism is a style which nicely describes the output generated via the AGP and has been used as an influence to further improve the system. The explicit removal of any natural or emotional content fits in with the constraints of auto-generated art and provides a good foundation for further investigation of aesthetic judgement. While the process may not be suitable for all scenarios, the benefits it does possess indicate that it is useful to constrain the available search space, improving the search for interesting items, in both automatic and human led systems.

However, the adherence to the Concretism style does introduce some potential downsides, it may limit the applicability of this system by producing content which is not aesthetically appealing to a wide audience. This, combined with limitations on the measures used in evaluation, may mean that the output is not judged to be aesthetically pleasing by a wide audience. The range of values across all the measures used indicate that this is likely not the case; however, additional evaluation is needed, where a human assessed understanding of the aesthetic appeal of the content is provided.

The AGP exposes multiple parameters improving the variety of items that can be generated and gives another dimension for the output to be explored. As the algorithm does not calculate a value for each pixel or voxel, the computational intensity of generating large or high-resolution items is reduced. This means that it is suitable for use in performance-intensive environments, such as Augmented and Virtual Reality as well as online systems. This opens up new possibilities for how the artwork can be utilised, such as creating continuously changing artwork which can be explored in Virtual Reality or just being able to generate a wide range of content suitable for display within a gallery.

One of the main themes of generating artwork using Evo-art systems is how the artwork can be assessed. As identified in the evaluation, existing formal measures are limited and will not fully represent human aesthetic judgement. This inferred disparity suggests that the contributing aspects of aesthetic judgement, and how these aspects interact with each other, when a judgement is being made, needs further investigation. The ability of the extended AGP to generate content with a wide range of values for these existing formal measures and due to the numerous ways the output from the AGP can be explored and experienced, means it is a good candidate to identify these further aesthetic aspects. This will help to map features of the artwork to high level aesthetic concepts, e.g., unfriendly which can be used by both automated systems and artists to more effectively generate artwork with these attributes. Ultimately, this allows a more intimate understanding of the aesthetic space to be achieved.

Declarations

Conflict of Interest The authors declare that they have no conflict of interest.

References

1. Acebo E, Sbert M. Benford's law for natural and synthetic images. In: First Eurographics conference on Computational Aesthetics in Graphics, Visualization and Imaging; 2005.
2. Bergen S, Ross BJ. Aesthetic 3d model evolution. Genet Program Evolvable Mach. 2013;14(3):339–67.
3. Birkhoff GD. Aesthetic measure. Mass: Cambridge; 1933.
4. Boden MA, et al. The creative mind: Myths and mechanisms. Psychology Press; 2004.
5. Byrne J, Hemberg E, O'Neill M, Brabazon A. A methodology for user directed search in evolutionary design. Genet Program Evolvable Mach. 2013;14(3):287–314.
6. Canaan R, Menzel S, Togelius J, Nealen A. Towards game-based metrics for computational co-creativity. In: 2018 IEEE conference on computational intelligence and games (CIG). IEEE, 2018; p. 1–8
7. Castelló P, Sbert M, Chover M, Feixas M. Techniques for computing viewpoint entropy of a 3d scene. In: International conference on computational science. Springer, 2006; p. 263–70
8. Cohen MW, Cherchiglia L, Costa R. Evolving mondrian-style artworks. In: International conference on evolutionary and biologically inspired music and art. Springer, 2017; p. 338–53.
9. Colton S. Automatic invention of fitness functions with application to scene generation. In: Workshops on applications of evolutionary computation. Springer, 2008; p. 381–91.
10. Colton S, Cook M, Raad A. Ludic considerations of tablet-based evo-art. In: European conference on the applications of evolutionary computation. Springer, 2011; p. 223–33.
11. Colton S. Evolving a library of artistic scene descriptors. In: International conference on evolutionary and biologically inspired music and art. Springer, 2012; p. 35–47.
12. Davies E, Tew P, Glowacki D, Smith J, Mitchell T. Evolving atomic aesthetics and dynamics. In: International conference on computational intelligence in music, sound, art and design. Springer; 2016. p. 17–30.
13. den Heijer E. Evolving art using measures for symmetry, compositional balance and liveliness. Int Conf Evolut Comput Theory Appl. 2012;2:52–61 (**SciTePress**).
14. den Heijer E, Eiben AE. Comparing aesthetic measures for evolutionary art. In: European conference on the applications of evolutionary computation. Springer, 2010; p. 311–20.
15. den Heijer E, Eiben A. Evolving pop art using scalable vector graphics. In: International conference on evolutionary and biologically inspired music and art. Springer, 2012; p. 48–59.
16. Easton E, Bernardet U, Ekart A. Tired of choosing? Just add structure and virtual reality. In: International conference on computational intelligence in music, sound, art and design (part of EvoStar). Springer; 2019. p. 142–55.
17. Easton E, Ekárt A, Bernardet U. Axial generation: a concretism-inspired method for synthesizing highly varied artworks. In: 10th international conference on artificial intelligence in music, sound, art and design, EvoMUSART 2021 held as Part of EvoStar 2021. Springer, 2021; p. 115–30
18. Easton E. Investigating user fatigue in evolutionary art. Master's thesis, Aston University; 2018.
19. Ecins A, Fermuller C, Aloimonos Y. Detecting reflectional symmetries in 3D data through symmetrical fitting. In: Proceedings of the IEEE international conference on computer vision workshops. 2017; p. 1779–83.

20. Ekárt A, Sharma D, Chalakov S. Modelling human preference in evolutionary art. In: European conference on the applications of evolutionary computation. Springer, 2011; p. 303–12.
21. Gircys M, Ross B. Image Evolution Using 2D Power Spectra. In: Complexity, vol. 2019. 2019. https://doi.org/10.1155/2019/7293193
22. Hollingsworth B, Schrum J. Infinite art gallery: a game world of interactively evolved artwork. In: 2019 IEEE congress on evolutionary computation (CEC), IEEE. 2019; p. 474–81.
23. Lehman J, Stanley KO. Exploiting open-endedness to solve problems through the search for novelty. In: ALIFE. 2008; p. 329–36.
24. Li Y, Hu C, Chen M, Hu J. Investigating aesthetic features to model human preference in evolutionary art. In: International conference on evolutionary and biologically inspired music and art. Springer, 2012; p. 153–64.
25. Machado P, Vinhas A, Correia J, Ekárt A. Evolving ambiguous images. AI Matters. 2015;2(1):7–8.
26. Machado P, Cardoso A. Computing aesthetics. In: Brazilian Symposium on Artificial Intelligence. Springer, 1998; p. 219–28.
27. Matkovic K, Neumann L, Neumann A, Psik T, Purgathofer W. Global contrast factor-a new approach to image contrast. Comput Aesthetics. 2005;2005:159–68.
28. McCormack J, Lomas A. Understanding aesthetic evaluation using deep learning. In: International conference on computational intelligence in music, sound, art and design (Part of EvoStar). Springer; 2020. p. 118–33.
29. McDermott J, Swafford JM, Hemberg M, Byrne J, Hemberg E, Fenton M, McNally C, Shotton E, O'Neill M. String-rewriting grammars for evolutionary architectural design. Environ Plann B Plann Des. 2012;39(4):713–31.
30. McDermott J. Graph grammars as a representation for interactive evolutionary 3D design. In: International conference on evolutionary and biologically inspired music and art. Springer, 2012; p. 199–210.
31. Mills A. Animating typescript using aesthetically evolved images. In: International conference on computational intelligence in music, sound, art and design. Springer; 2016. p. 126–34.
32. Muehlbauer M, Burry J, Song A. Automated shape design by grammatical evolution. In: International conference on evolutionary and biologically inspired music and art. Springer, 2017; p. 217–29.
33. Nguyen AM, Yosinski J, Clune J. Innovation engines: automated creativity and improved stochastic optimization via deep learning. In: Proceedings of the 2015 annual conference on genetic and evolutionary computation. 2015; p. 959–66.
34. Nicolau M, Costelloe D. Using grammatical evolution to parameterise interactive 3D image generation. In: European conference on the applications of evolutionary computation. Springer, 2011; p. 374–83.
35. O'Neill M, McDermott J, Swafford JM, Byrne J, Hemberg E, Brabazon A, Shotton E, McNally C, Hemberg M. Evolutionary design using grammatical evolution and shape grammars: designing a shelter. Int J Design Eng. 2010;3(1):4–24.
36. O'Reilly UM, Hemberg M. Integrating generative growth and evolutionary computation for form exploration. Genet Program Evolvable Mach. 2007;8(2):163–86.
37. Rigau J, Feixas M, Sbert M. Conceptualizing Birkhoff's aesthetic measure using Shannon entropy and Kolmogorov complexity. In: Computational Aesthetics. 2007; p. 105–12.
38. Rigau J, Feixas M, Sbert M. Conceptualizing Birkhoff's aesthetic measure using Shannon entropy and Kolmogorov complexity. In: Computational Aesthetics. 2007; p. 105–12.
39. Secretan J, Beato N, D Ambrosio DB, Rodriguez A, Campbell A, Stanley KO. Picbreeder: evolving pictures collaboratively online. In: Proceedings of the SIGCHI conference on human factors in computing systems. 2008; p. 1759–68.
40. Sims K. Artificial evolution for computer graphics. In: Proceedings of the 18th annual conference on computer graphics and interactive techniques. 1991; p. 319–28.
41. Takagi H. Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation. Proc IEEE. 2001;89(9):1275–96.
42. Tate Concrete art. 2017. https://www.tate.org.uk/art/art-terms/c/concrete-art/, [Accessed: 2020-11-20].
43. Tinio PP, Leder H. Just how stable are stable aesthetic features? Symmetry, complexity, and the jaws of massive familiarization. Acta Physiol (Oxf). 2009;130(3):241–50.
44. Tweraser I, Gillespie LE, Schrum J. Querying across time to interactively evolve animations. In: Proceedings of the genetic and evolutionary computation conference. 2018; p. 213–20.
45. Helion, J. Art Concret. 1930. [online] Available at: https://monoskop.org/images/2/2d/Art_concret_1_1930.pdf. Accessed 20 Nov 2020.
46. Vinhas A, Assunção F, Correia J, Ekárt A, Machado P. Fitness and novelty in evolutionary art. In: International conference on computational intelligence in music, sound, art and design. Springer; 2016. p. 225–40.
47. Wiggins GA. A preliminary framework for description, analysis and comparison of creative systems. Knowl-Based Syst. 2006;19(7):449–58.
48. Zeki S. Inner vision: An exploration of art and the brain. J Aesthet Art Crit 2002;60(4).