Exploring Human Models of Innovation for Generative AI

Gualtiero B. Colombo, Hantao Liu, Roger M. Whitaker
School of Computer Science and Informatics
Cardiff University
Cardiff, CF24 4AG UK
ColomboG@cardiff.ac.uk, LiuH35@cardiff.ac.uk, whitakerrm@cardiff.ac.uk

Abstract

The ability to innovate is a precious commodity that humans are well-disposed to accomplishing. Currently, in the quest for development of artificial intelligence that is generative, blueprints for innovation are important for consideration. However those following principles of human innovation have been somewhat overlooked. The field of cultural evolution presents interesting models for explaining human innovation as evolutionary processes driven by intelligent biases. These offer approaches that can be followed in an algorithmic form by the machine, with particular degrees of freedom concerning bias. In this paper we take a first step in exploring cultural evolution models for generative AI. In particular, we use the concepts of cultural selection and biased transformation, where changes are driven by the bias imparted at different points of human decision making processes, in addition to social learning from others. We develop these approaches using a population of neural networks, each capable of drawing an image. We explore how neural networks and the resultant art evolve under these alternative interpretations of cultural evolution, using biases based on preferred prior images. The investigation suggests that bias must also evolve for innovation to persist, something which is given little emphasis in the literature.

Introduction

Interest in enabling computers to support creativity is something that is as old as computer science itself. Indeed, in an essay dating back to 1948 (Turing 1948), Alan Turing, a father of modern computing, speculated that different forms of search would need to support "intelligent machinery", and highlighted what he called "cultural search" as a process aligned to the mission of human creativity (Turing 1950). In very recent times we are now seeing considerable progress in artificial intelligence (AI) that is creative in open-ended scenarios, particularly through *generative AI*. Applications such as ChatGPT (Schulman et al. 2022), DALL-E 2 (Ramesh et al. 2022), and Stable Diffusion (Rombach et al. 2022) can create meaningful and complex content in response to limited instructions, for scenarios of complexity beyond which we have previously seen (Oppenlaender 2022). This is particularly the case for generative art and text, where the underlying models are highly dependent on scale (Galanter 2016; Boden and Edmonds 2009). For example, the training of large language models requires billions of words, bringing into question whether as an alternative, there are useful general principles through which meaningful creativity can established by alternative computational processes (Floridi and Chiriatti 2020; Dale 2021; Berns et al. 2021). To this end, current successful approaches include generative adversarial networks, where creativity is driven by competition (Tan et al. 2017), racheting training between a generating and discriminating neural network that results in high quality generation for open-ended problems including art (Shahriar 2022).

Less well considered are the processes underpinning the evolution of human innovation. The rationale for these models is strong, as humans have been able to innovate beyond all other species. Fundamental insights on the nature of creative processes have spanned both computing and psychology, with contributions such as those from Boden and Sawyer (Boden 2005; 2004; Sawyer 2011) indicating methods through which the individual mind can achieve creativity. These contributions serve to counter the illusion that creativity is a form of "magic", instead being methods that allow large search spaces to be navigated. It is also worth noting the often highly social nature of creativity, in that it is rarely achieved in isolation, and is progressively developed by building on the achievements of others, or from "*standing on the shoulders of giants*". This is something that Turing articulated in his early treatment of this subject (Turing 1948).

Today this area is recognised as *cultural evolution* (Tomasello 2009; Boyd and Richerson 1988; Mesoudi 2011), a cross-disciplinary endeavour that broadly seeks to understand how human innovations take hold (Tomasello, Kruger, and Ratner 1993). Here innovation has a specific meaning, representing the combination of *invention* and *social learning* (Paulus and Dzindolet 2008). In this context creativity is the process supporting invention, and this may influenced by others. Through cultural evolution, *cumulative culture* (Mesoudi and Thornton 2018) is now seen as a front runner in explaining how humans have become supreme innovators as compared to all other species. The main premise of cumulative cultural evolution is the concept of *ratcheting*, where improvements and novelty build without reverting to previous states over the longer term (Tennie, Call, and Tomasello 2009). This allows creativity to build in sophistication. It is argued that this presents similarities to Darwinian evolution (Mesoudi 2011), albeit with different mechanisms that allow change to happen much more quickly. Although there is significant debate about how cumulative culture results, at least two key models have emerged, known as *biased transformation* and *cultural selection* (Mesoudi 2021). These models expose the crucial role that human bias plays in executing cultural evolution, and can be approximated as simple algorithms that offer degrees of freedom as to their interpretation, configuration and sophistication. They can also function without necessarily pre-training, using association with memories to impart biases, which is often aligned with human decision making.

Contribution

Our overall interest is in techniques that are able to *persistently create artifacts that increasingly embody innovation and novelty*. Based on the success of humans in achieving this, in this paper we focus on two key mechanisms of cultural evolution, namely *cultural selection* and *biased transformation*. We adopt them as an inspiration for computational techniques where a small group of retained preferences are collectively engaged in creating new art represented through neural networks. Note that art is chosen as a vehicle to study creativity because it is readily accessible for human interpretation of innovation.

In general, innovation can be a challenging concept to measure because features that are innovative may not be foreseen, impeding quantitative measures. We use a general approach where instances of art are each drawn by a neural network and a set (or population) of images are maintained. Bias through images (represented by neural networks) are used to effectively represent preferences and a persistent memory, referred to as *preferred priors*. We explore ways in which bias can be applied through neural networks that are engaged in creating art, and we include the use of techniques from neuro-evolution to combine and impart bias. This approach allows directed modifications to be explored, where new images are created that are then available to repeatedly build upon in future. Since images are easily human interpretable we are able to gain a first-hand qualitative understanding of the role of bias in creative computational processes aligned to cultural evolution. This allows us to develop new insights and hypotheses.

The Neural Network Artist

Artificial neural networks (ANNs) are a mainstay of current AI, being robust to scaling and applicable across wide ranging scenarios. We focus on a novel form of ANNs called Compositional Pattern Producing Networks (CPPNs), that can be used to create images (Stanley 2007). Interesting images can be produced even from simple CPPN structures - an example is shown in Figure 1. CPPNs function by taking the x and y coordinate of an image's pixel as the input, and return as output the colour of the pixel. Thus each CPPN can be thought of as an artist that has created a painting. The structure of the CPPN and the activation functions used on the CPPN's nodes determine the form of the image that is produced. The previous use of CPPNs for image

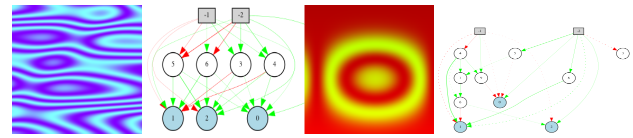

random CPPN Image e CPPN

Figure 1: CPPNs representing a simple randomly generated image and a more complex 'prior image'.

generation has been highly successful, for example being used in participatory crowdsourcing experiments (Secretan et al. 2008) that transformed abstract art to meaningful or interesting compositions. Using CPPNs provides a way to expose the "brain" of the virtual artist, and a novel point of influence. For example it is possible to make small mutations to a CPPN, such as a random change to the function of a node or the strength of connection between nodes (Stanley 2007), resulting in perturbations to the associated image. More directed changes can be invoked through combining CPPNs so that characteristics from one image can be used to influence another. This is non-trivial because the structure of CPPNs may vary making it challenging to map between such neural networks. However techniques from neuro-evolution can be used to create this integration, such as the crossover mechanism used in the NEAT algorithm (Stanley and Miikkulainen 2002; Stanley, D'Ambrosio, and Gauci 2009). Consequently, CPPNs and the approach to combining them through crossover give a means to impart bias in a cultural evolutionary process. These capabilities allow new computational models based on cultural evolution to be explored. In particular, it is interesting to see the extent to which they can be harnessed to generate novel artifacts with persistent creativity and increased complexity. This aspect gives an important insight concerning open-endedness, which is the grand-challenge of creating an algorithm that persistently creates innovations (Stanley 2019; Lehman, Stanley, and others 2008; Stanley, Lehman, and Soros 2017).

Experimental approach

To explore models of cultural evolution in this context, we address two directions for experimentation. Firstly, to understand the search space and the impact of bias, we consider biased and non-biased navigation through the search space, and the differences between them (Experiment 1). This requires consideration of the effects on changes to evolving a selected starting image. Secondly, to understand cultural models, we consider how approximations to biased transformation and cultural selection perform (Experiment 2), using retained memories of preferred prior images as the basis for bias. Note that necessarily, the evaluation here is subjective and exploratory, as a necessary first step to be further developed. The benefits of using art are in the human interpretability of the artifacts and their novelty, but this can only be judged by selective qualitative means. Nevertheless, this provides a useful basis to develop observations and hypotheses for more rigorous future investigation.

Experiment 1: Mutating a Single Image

What happens when we mutate a CPPN, or more specifically, how is the resulting image disrupted? This is a fundamental question underlying models of cultural evolution. Changes to artifacts, whether intentional or random, are the basis for *variation*. Variation underpins all evolutionary processes, without which artifacts remain in stasis. Variation opens up choice for further modification, and it is the basis for further random or directed change. The choices and modifications that are made determine how creativity emerges. While the human can impart intuition in selection and transformation decisions, using embedded skills, experiences and memories, the computer requires explicit bias to be programmed.

Algorithmic Approach

To assess this we compare two alternative strategies. Firstly taking a single starting point CPPN, denoted I, that has been randomly generated (and thus maps to a random image), we apply small random changes to the edge weighting or the activation function within a node from a hidden layer of the CPPN. This is successively repeated to the CPPNs that result, being equivalent to a random walk through the search space, representing an unguided and uninfluenced path. The general approach used is presented in Algorithm 1. This involves starting with a CPPN I (line 2) and evolving this using a simple procedure where n_s alternative random mutations are made to I (called the candidate set - line 7) from which one is randomly selected to update I (lines 8 and 9). This approach to randomisation is used so that it exploits a common framework of code. The output from Algorithm 1 gives a baseline to understand characteristics of a non-biased random approach.

Secondly, we consider biased transformation, where directed search is applied (Algorithm 2) in place of the small random changes considered in Algorithm 1. Algorithm 2 has two variations embedded within it, based on the strength of directed transformation. Initially a set of preferred prior images, denoted P, are defined (line 1). These provide the source of bias - in effect memories of interesting images represented by CPPNs. The starting CPPN for subsequent evolution is initalised in line 2. Note that this could also be one of the preferred priors.

The first step is to create a set of alternatives from I, which is called the *candidate set*. This is populated in two alternative ways (line 7), either mutations of I (called $biasedTrans_1$) or a crossover and mutation between I and a random CPPN from the set of priors (called $biasedTrans_2$). Crossover is a mechanism that combines two CPPNs and imparts characteristics from both to form a new CPPN, akin to creating an offspring. It is a non-trivial operation because CPPNs are not guaranteed to have the same structure. Here we use the crossover operator defined in (Stanley and Miikkulainen 2002). The resulting CPPN I' is subject to a small mutation (line 8) to guard against I or a prior being over-represented in the candidate set. Finally, a selection is made from the candidate set and this involves bias aligned to a particular prior. To achieve this, at each iteration, a preferred prior of interest is randomly chosen from P, denoted P_i (line 5). Then a subset of the candidate set is identified, involving a selection of most similar n_i images to P_i. This subset is called the *individuals set* (line 10). This excludes those images that are less well-related to P_i. Note that similarity is applied to the images rather than the CPPNs from which they are defined. This requires techniques from image processing and we employ a deep residual neural network Resnet (He et al. 2016), trained on the over 20 million images from the Imagenet dataset, further fine tuned with the methodology presented in (Wang et al. 2014). This allows general abstract features of images to drive similarity, such as shapes and patterns. It is ideal for our needs because it ensures similarity isn't based on over-fitting, being in more keeping with human intuition rather than precision. From the individuals set, then finally a random selection is made (line 11), which is updates I.

For exploratory purposes, a candidate set with cardinality 50 has been used, alongside an individuals set of size 10 and we run the algorithms for 50 iterations. Mutation settings for the CPPNs are set as equal to the default from (McIntyre et al. 2015).

Algorithm 1 Random Walk

1: **procedure** RANDOM WALK(Evolving CPPN I, Candidate Set Size n_s, Number of Iterations n^i)
2: $I \leftarrow P_j \in P$ **or** $I \leftarrow RandomImage$; $i = 0$
3: **while** $i < n^i$ **do**
4: CandidateSet = $\emptyset$; $j = 0$
5: **for** $j < n_s$ **do**
6: $I' = mutate(I)$
7: CandidateSet $\leftarrow$ CandidateSet $\cup$ I'; j++;
8: $V' \leftarrow$ Select Randomly from CandidateSet
9: $I \leftarrow V'$; i++;

Algorithm 2 $biasedTrans_1$ and $biasedTrans_2$

1: **procedure** BIASEDTRANSFORMATION(Set of Prior Images P, Evolving CPPN I, CandidateSetSize n_s, IndividualsSetSize n_i, NumberOfIterations n^i)
2: $I \leftarrow P_j \in P$ **or** $I \leftarrow RandomImage$; $i = 0$
3: **while** $i < n^i$ **do**
4: CandidateSet = $\emptyset$, IndividualsSet = $\emptyset$; $j = 0$
5: Set Current Prior as $P_i \in P$
6: **for** $j < n_s$ **do**
7: $I' \leftarrow I$ **or** $\triangleright$ $biasedTrans_1$
$I' \leftarrow$ Crossover(I, P_i) $\triangleright$ $biasedTrans_2$
8: $I'' = mutate(I')$
9: CandidateSet $\leftarrow$ CandidateSet $\cup$ I''; j++;
10: Select IndividualsSet $\subset$ CandidateSet as the Set of n_i Images in CandidateSet Most Similar to Current Prior P_i
11: $V' \leftarrow$ Select Randomly from IndividualsSet
12: $I \leftarrow V'$; i++;

Experiment 1: Results

Firstly we consider the effects of random walk through the search space, based on Algorithm 1. To demonstrate

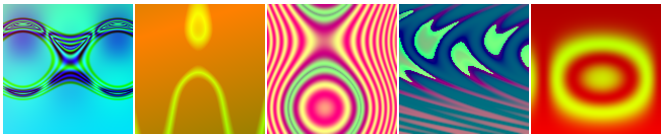

Image a Image b Image c Image d Image e

Figure 2: Example of five images represented through CPPNs that are used as candidates preferred priors and/or starting points in various experiments. These are chosen for having little similarity between them based on ResNet assessment.

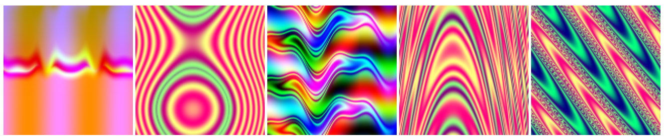

Image f Image g Image h Image i Image j

Figure 3: Example of five images represented through CPPNs that are used as candidates preferred priors in various experiments. These are chosen for having greater similarity between them based on ResNet assessment.

this, five sample images from Figure 2 and their associated CPPNs were taken as starting points, and the random evolution of resulting images were observed. Figure 4 shows interesting snapshots from the resulting sequence of images produced. Relatively quick changes occur across the images, with particularly interesting iterations highlighted in Figure 4, where the relationships between images can be observed. As expected, there are no overall patterns that can observed. Some random paths become complex in different ways (e.g., images a and d) while others remain similarly complex (e.g., image b), although they all tend to become more abstract and less well aligned to particular shapes that humans can identify with. Others drift to lower complexity (e.g., c and e). What is evident though is the rich variety of images that can be readily generated and how easy it is to transverse the search space.

In contrast to a random walk, the results from biased transformation (Algorithm 2) are stark, where the impact of biased transformation is significant. Figures 5 and 6 demonstrate this using a common starting CPPN (see Figures 1 and 2), and they present selections from some of the most interesting images created across different runs with alternative random seeds. These are of course subjective selections, but are representative of the search space. In each of Figures 5, 6, 7 and 8 different priors are applied, ranging from two to five priors. Also two alternative starting images are used, to provide a further comparison. Finally we also consider difference in applying a set of highly similar priors (Figure 3) as compared to a set that are mutually dissimilar, based on ResNet similarity (Figure 9).

Using qualitative inspection from experimentation, we draw the following observations and hypotheses. Firstly, $biasedTrans_2$ has a much stronger evolutionary impact in terms of diversity of interesting images than for $biasedTrans_1$. In other words, introducing crossover provides a strong influence to direct the evolution towards creative areas of the search space related to preferred priors. Secondly, although a minimal number of preferred priors will support the discovery of interesting solutions, additional priors appear to increase the diversity and complexity of the most interesting images that are discovered. This is especially the case when crossover is involved ($biasedTrans_2$). However, fully displaying this using a limited number of images is challenging. Thirdly, in all cases, using alternative random seeds is sufficient to drive the evolution in diverse directions across the search space. This seemed to be amplified when a greater number of priors were involved, or when crossover was employed ($biasedTrans_2$). Fourth, it is evident that both similar and dissimilar sets of prior images can promote creativity, while the dissimilar priors seem to influence shape formation, while similar priors seem to add more intricate detail to images. Finally, the results seem to affirm that the concept of similarity to retained memories (or preferred priors) is sufficient to drive creativity within a few iterations without new images being directly tied to the priors from which they have been influenced.

Experiment 2: Cultural Evolution on a Population of Images

To further explore models of cultural evolution, it is necessary to introduce a population of alternative artifacts that are available to be updated, providing a diversity through which innovations can accumulate and transfer between individuals. This is a basis for two key models of cultural evolution, namely *cultural selection* and *biased transformation* (Mesoudi 2021). These models allow external influence from multiple sources each represented by a *preferred prior* image to possibly influence decision making when acting upon a population of artifacts of some description. The difference between cultural selection and biased transformation concerns where bias comes into play - either at the point of selection of artifact to modify (cultural selection) or the way in which they are modified (biased transformation). In reality both these elements may combine (Mesoudi 2021) but it is prudent, from an exploratory perspective, to understand the differences between alternative models.

Algorithmic Approach

We use the algorithmic framework analogous to that described in the previous section and focus on evolving a population of random CPPNs aligned to different models of cultural evolution. Firstly we focus on the effects of bias at the selection stage, aligned to *cultural selection*. In Algorithm 3 we evolve a population of N CPPNs (set Pop) by replacing at each generation a subset of Pop, denoted *ReplacedSet*, and replacing it with another set of CPPNs, denoted *NewIndSet*.

To decide how *NewIndSet* is composed, we firstly create a set, *CandidateSet* (lines 8-11) from the population. This can be created in two alternative ways (line 9), either through mutating a randomly selected individual (which we

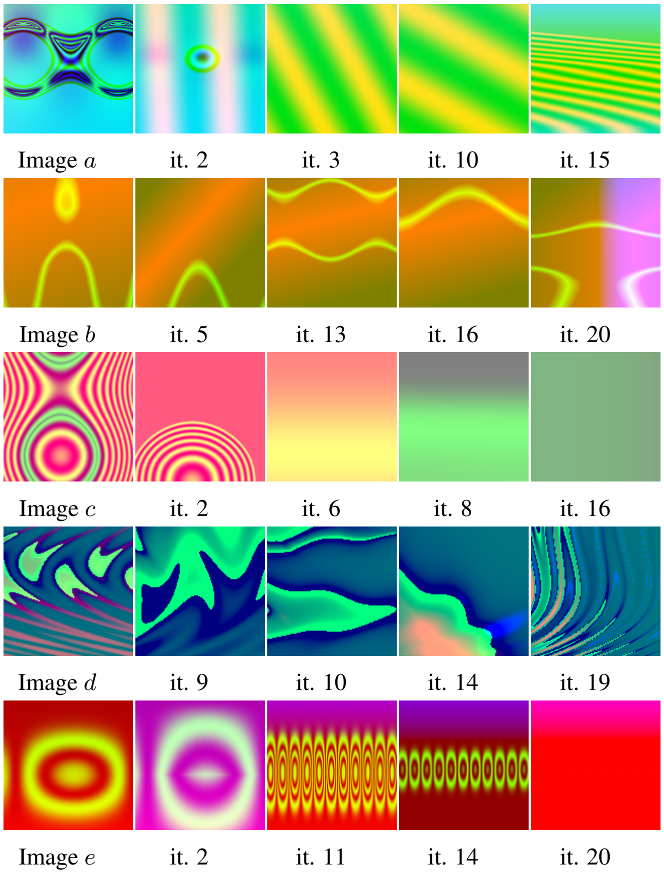

Figure 4: Examples of representative and interesting images produced under Algorithm 1 (random walk). Each row represents a sample of images selected from random evolution from five starting points (Images a - e). Numbers under images indicate the iteration from which they are taken.

call *random selection*), or by making the selection based on weighted similarity with a randomly selected prior P_i (line 6), representing a form of direct bias (Boyd and Richerson 1988), and then applying a mutation. We call this approach *cultural selection*. The *NewIndSet* is then randomly selected as a subset of *CandidateSet*. Finally, the set of images that are least similar to P_i are then removed from the population *Pop* (*ReplacedSet* in lines 16-18). This biased removal is important in various approaches to cultural evolution (Boyd and Richerson 1988; Dawkins and others 1996).

Algorithm 4 evolves a population of CPPNs based on *biased transformation*, where the transformation of artifacts from the population are subject to biases. This contrasts to Algorithm 3 where the bias is applied to selection. As in Algorithm 3, at each generation Algorithm 4 takes a subset of Pop, denoted *ReplacedSet*, and replaces it with another set of CPPNs, denoted *NewIndSet*. Members of the *CandidateSet* are initially selected at random (line 9). These are then either directly mutated under $biasedTransformation_1$ or crossover occurs with P_i, representing a form of guided variation (Boyd and Rich-

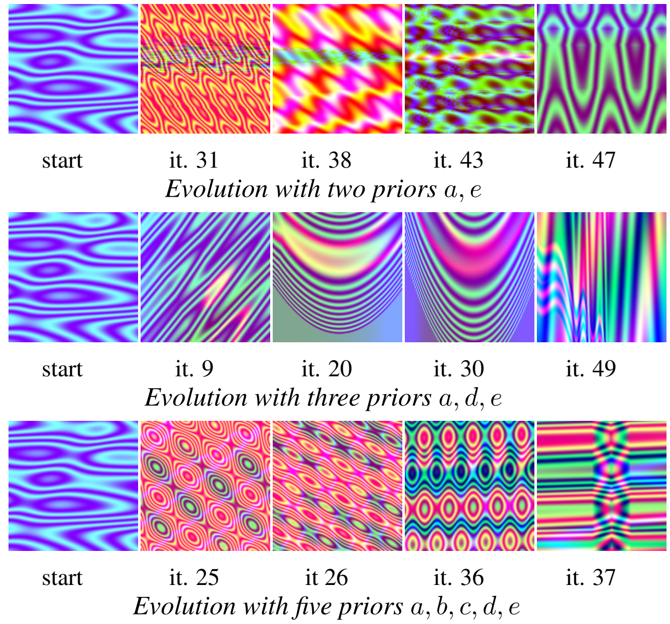

Figure 5: Examples of the most creative and interesting images produced under $biasedTrans1$ (Algorithm 2) for two (top), three (middle), and five (bottom) prior images using a randomly generated image as starting point. Numbers under images indicate the iteration from which they are taken.

erson 1988), denoted $biasedTransformation_2$ (lines 10-

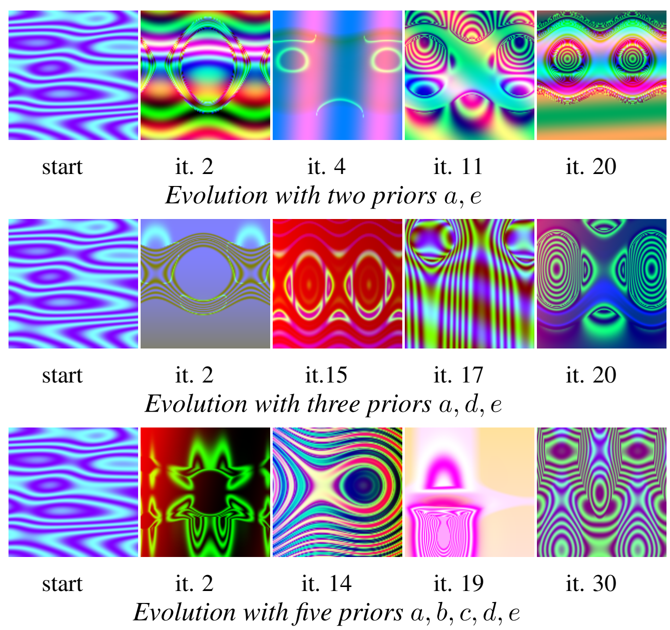

Figure 6: Examples of the most creative and interesting images produced under $biasedTrans2$ (Algorithm 2) for two (top), three (middle), and five (bottom) prior images using a randomly generated image as starting point. Numbers under images indicate the iteration from which they are taken.

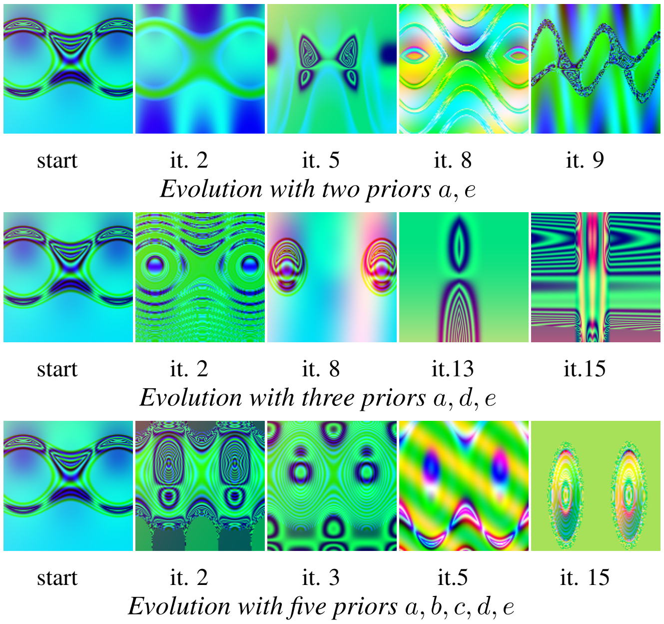

Figure 7: Examples of the most creative and interesting images produced under $biasedTrans1$ (Algorithm 2) for two (top), three (middle), and five (bottom) prior images using Image a as starting point. Numbers under images indicate the iteration from which they are taken.

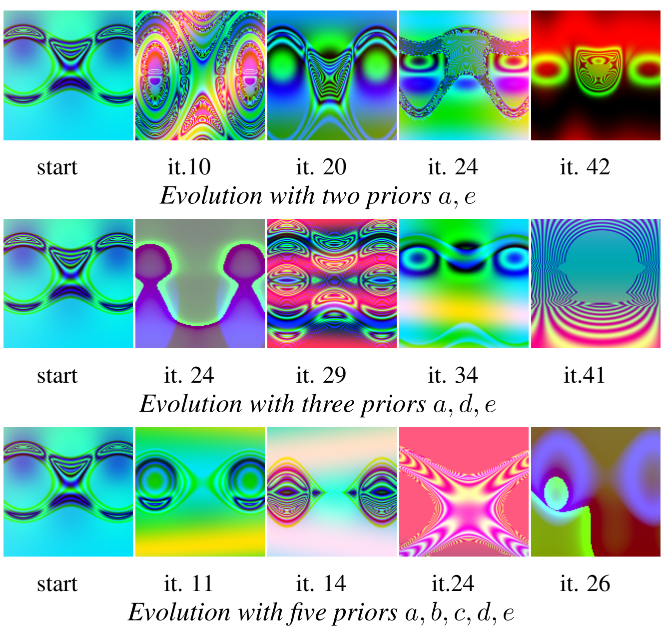

Figure 8: Examples of the most creative and interesting images produced under $biasedTrans2$ (Algorithm 2) for two (top), three (middle), and five (bottom) prior images using Image a as starting point. Numbers under images indicate the iteration from which they are taken.

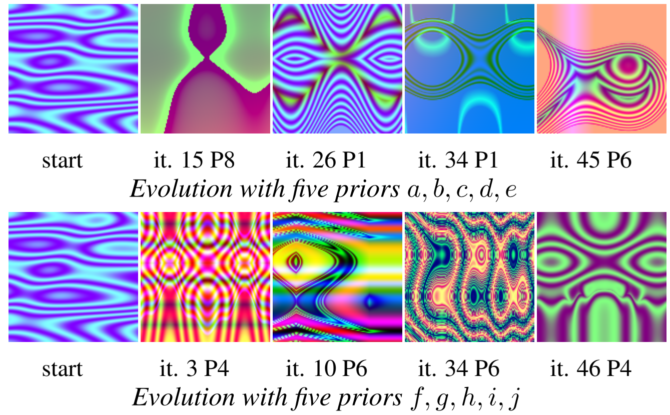

Figure 9: Examples of the most creative and interesting images produced under $biasedTrans2$ (Algorithm 2) for similar (top) and dissimilar (bottom) sets of priors using a randomly generated image as starting point. Numbers under images indicate the iteration from which they are taken.

11). From *CandidateSet*, we then select a subset *NewIndSet* that is most similar to P_i (lines 13-15). Finally, the set of images that are least similar to P_i are then removed from *Pop* (*ReplacedSet* in lines 16-18) and *NewIndSet* is added to *Pop*. This replacement approach follows that in Algorithm 3. For exploratory purposes, a candidate set with cardinality 250 has been used, alongside an individuals set of size 10, a replaced set also of size 10, and we run the algorithms for 50 generations.

Experiment 2: Results

Firstly we consider evolving a population of CPPNs ($N = 50$) and adopt only one preferred prior, taking image e from Figure 3. This is designed to test the potential convergence differences between the variations in Algorithms 3 and 4. The cultural evolution literature (Mesoudi 2011; Boyd and Richerson 1988) indicates that while both culture selection and biased transformation support convergence in simple fitness-based models, biased transformation is quicker.

The results from Figure 10 show that biased transformation has a significant effect in directing the evolution based on the preferred prior, as compared to both cultural and random selection. This is particularly the case for $biasedTransformation_2$. Similarity here is measured using ResNet-based measure (Wang et al. 2014). Note that all approaches involve biased removal of CCPNs from the population at the point at which replacements are added. The results also show that there are fundamental limitations in applying only a single point of bias, in this case a single preferred prior, because evolution gets drawn towards the original point of bias, as seen in Figure 11. This is also noted in cultural evolution treatments of biased transformation.

Secondly, based on using a set of five preferred priors as bias, we consider the characteristics of cultural evolution based on Algorithms 3 and 4. In Figure 14 we present a selection of images representative of the most novel and creative images from across different generations for each of the four techniques presented through Algorithms 3 and 4.

Algorithm 3 Random Selection and Cultural Selection

procedure RANDOMSELECTION / CULTURAL SELECTION (NumberofGens G, EvolvingPupulationOfImages Pop, SetofPriorImages P, PopSize N, CandidateSetSize n_s, NewIndSetSize n_i, ReplacedSetSize n_r)
  $NewIndSet{=}\emptyset, CandidateSet{=}\emptyset, ReplacedSet{=}\emptyset$
  Set Pop as a Population of Random Images; $g = 0$
  while g$<G$ **do**
    Select Randomly Current Prior P_i from P
    $i = 0;\ j = 0;\ k = 0$
    for $i < n_s$ **do**
      $I \leftarrow$ Select Randomly from Pop ▷ rand. sel.
      or
      $I \leftarrow$ Select from Pop Proportionally to Similarity to P ▷ cult. sel.
      $I' \leftarrow mutate(I)$
      $CandidateSet \leftarrow CandidateSet \cup I'; i{+}{+}$
    for $j < n_i$ **do**
      $V \leftarrow$ Select Randomly from CandidateSet
      $NewIndSet \leftarrow NewIndSet \cup V; j{+}{+}$
    for $k < n_r$ **do**
      $K \leftarrow$ Set of $ReplacedSetSize$ Images in Pop Least Similar to P_i
      $ReplacedSet \leftarrow ReplacedSet \cup K; k{+}{+}$
    $Pop \leftarrow Pop - ReplacedSet$
    $Pop \leftarrow Pop \cup NewIndSet; g{+}{+}$

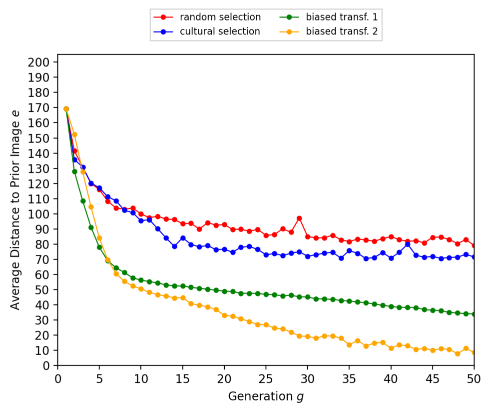

Figure 10: Average distance of population to image e as the preferred prior produced from the variations in Algorithms 3 and 4.

From inspection and multiple trials, greater levels of creativity seems apparent under biased transformation. Aligned to these experiments, we also track the average similarity in the population of CPPNs as compared to each of the five priors. This is presented for *cultural selection* (Figure 12) and $biasedTransformation_2$ (Figure 13). We note that there are considerable differences - while Figure 12 exhibits general trends towards similarity and convergence, the alternative is true in Figure 13.

Based on our qualitative observation from experimentation we draw the following observations and hypotheses. Firstly, although bias can drive creativity, it can also equally restrict creativity and innovation if biases are restricted, as seen when only a single prior is applied. Thus multiple sources of bias have an important effect on the creativity that is achieved. Secondly, multiple sources of bias, combined with the evolution of a population of images rather than evolution of a single image (Experiment 1) provides the opportunity for much diversity of images to emerge. This is more strongly felt under biased transformation, when the transformation is heavily directed through crossover-based techniques ($biasedTransformation_2$). Finally, from observing Figures 12 and 13 we hypothesise that tensions between biases can drive creativity and the innovation that results. In effect biases, when heavily directed through biased transformation, are steering a path through interesting elements of the search space, seemingly allowing more scope to evolve shapes for example, rather than only adding complexity to an existing form of image.

Algorithm 4 $biasedTrans_1$ and $biasedTrans_2$

procedure BIASEDTRANSFORMATION (NumberofGens G, EvolvingPupulationOfImages Pop, SetofPriorImages P, PopSize $= N$, CandidateSetSize n_s, NewIndSetSize n_i, ReplacedSetSize n_r)
  $NewIndSet{=}\emptyset, CandidateSet{=}\emptyset, ReplacedSet{=}\emptyset$
  Set Pop as Population of Random images; $g = 0$
  while g$<G$ **do**
    Select Randomly Current Prior P_i from P
    $i = 0;\ j = 0;\ k = 0$
    for $i < n_s$ **do**
      $I \leftarrow$ Select Randomly from Pop
      $I' \leftarrow I$ ▷ $biasedTrans_1$
      or
      $I' \leftarrow crossover(I, P_i)$ ▷ $biasedTrans_2$
      $I'' = mutate(I')$
      $CandidateSet \leftarrow CandidateSet \cup I'; i{+}{+}$
    for $j < n_i$ **do**
      $V \leftarrow$ Set of $NewIndSetSize$ Images in $CandidateSet$ Most Similar to P_i
      $NewIndSet \leftarrow NewIndSet \cup V; j{+}{+}$
    for $k < n_r$ **do**
      $K \leftarrow$ Set of $ReplacedSetSize$ Images in Pop Least Similar to P_i
      $ReplacedSet \leftarrow ReplacedSet \cup K; k{+}{+}$
    $Pop \leftarrow Pop - ReplacedSet$
    $Pop \leftarrow Pop \cup NewIndSet; g{+}{+}$

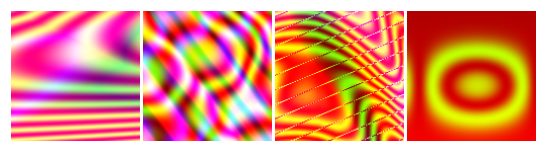

random selection | cultural selection | biased trans. 1 | biased trans. 2

Figure 11: Images with greatest similarity to image e as the preferred prior produced from the variations of Algorithms 3 and 4 over 50 generations.

Conclusion and Future Work

This study has brought together diverse techniques from neural networks, neuro-evolution and visual computing to support a new exploratory approach for harnessing computational creativity. These techniques have been used to explore whether fundamental human models of innovation, known as *cultural evolution*, can inspire new computational techniques to generate creativity from minimal user input and computational forms of bias. Models of cultural evolution represent important techniques because they capture the ways in which humans have been supremely successful as innovators. Art provides an excellent vehicle for exploring computational techniques in this context, with neural networks creating images for human interpretation of novelty. Our initial findings show prospects for new computational techniques based on cultural evolution. A key issue concerns the role of bias, and *biased transformation* in particular, shows promise. Imparting bias in a computational form can be a challenge, and the approach undertaken here supports the idea that retained memories, and abstract similarity to them, can function as an effective method. In other words, ideas from the past, or embedded preferences, can shape the creation of new artifacts in unforeseen and novel ways.

To understand the impact of biases we explored biased and non-biased navigation through the search space. Furthermore, we considered how approximations to biased transformation and cultural selection perform, to understand how these cultural models impact on algorithmic creation. Our investigation highlights the neural networks' ability to create images that evolve based on biased transformation, suggesting that further development should focus on the issue of bias and also that bias must evolve for innovation to persist. Novel forms of learning can be considered to allow adaption of bias in step with the level of complexity in the population. Little emphasis is given to the dynamics of bias in the cultural evolution literature, which often features snapshots of dynamic behaviour or static fitness functions as a proxy for bias. However machine learning and computational evolutionary techniques offer new prospects for achieving this. We believe that this could be an important aspect in developing persistent innovation aligned to open-endedness (Stanley, Lehman, and Soros 2017).

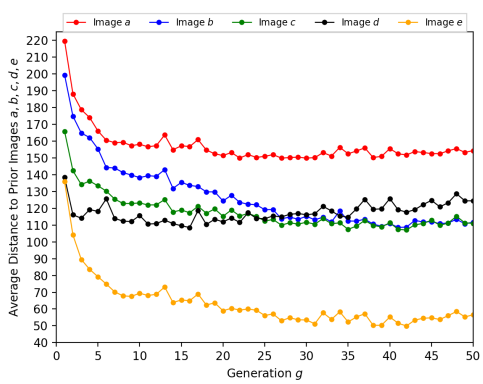

Figure 12: Average distance of population to the each of five prior images a, b, c, d, e over generations under *cultural selection* (Algortihm 3).

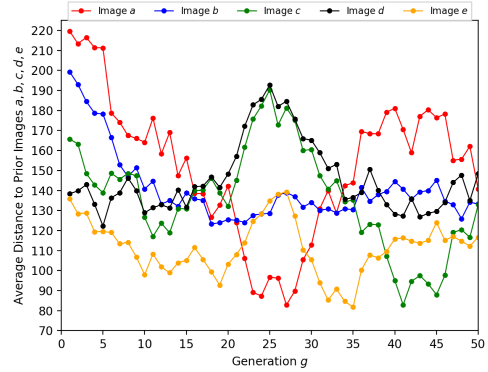

Figure 13: Average distance of population to each of the five prior images a, b, c, d, e over generations under $biasedTrans2$ (Algorithm 4).

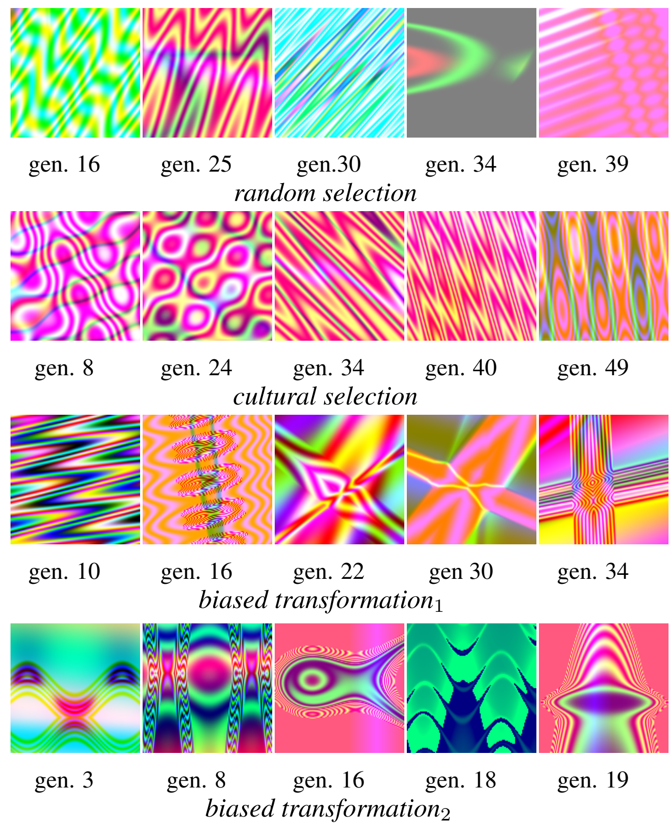

Figure 14: Examples of the most creative and interesting images from the population produced under the variations in Algorithms 3 and 4. A population of random images was used as starting point. Numbers under images indicate the generation from which they are taken.

Author Contributions

All authors contributed to the model, experimental design and analysis. Author 1 led the implementation. Author 2 provided additioanl expertise concerning visual computing. All authors contributed to drafting and reviewing the manuscript.

References

Berns, S.; Broad, T.; Guckelsberger, C.; and Colton, S. 2021. Automating generative deep learning for artistic purposes: Challenges and opportunities. *arXiv preprint arXiv:2107.01858*.

Boden, M. A., and Edmonds, E. A. 2009. What is generative art? *Digital Creativity* 20(1-2):21–46.

Boden, M. A. 2004. *The creative mind: Myths and mechanisms*. Routledge.

Boden, M. A. 2005. What is creativity? In *Creativity in human evolution and prehistory*. Routledge. 27–55.

Boyd, R., and Richerson, P. J. 1988. *Culture and the evolutionary process*. University of Chicago press.

Dale, R. 2021. Gpt-3: What's it good for? *Natural Language Engineering* 27(1):113–118.

Dawkins, R., et al. 1996. *The blind watchmaker: Why the evidence of evolution reveals a universe without design*. WW Norton & Company.

Floridi, L., and Chiriatti, M. 2020. Gpt-3: Its nature, scope, limits, and consequences. *Minds and Machines* 30:681–694.

Galanter, P. 2016. Generative art theory. *A companion to digital art* 146–180.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

Lehman, J.; Stanley, K. O.; et al. 2008. Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*, 329–336.

McIntyre, A.; Kallada, M.; Miguel, C. G.; Feher de Silva, C.; and Netto, M. L. 2015. neat-python.

Mesoudi, A., and Thornton, A. 2018. What is cumulative cultural evolution? *Proceedings of the Royal Society B* 285(1880):20180712.

Mesoudi, A. 2011. *Cultural evolution - How Darwinian Theory Can Explain Human Culture and Synthesize the Social Sciences*. University of Chicago Press.

Mesoudi, A. 2021. Cultural selection and biased transformation: two dynamics of cultural evolution. *Philosophical Transactions of the Royal Society B* 376(1828):20200053.

Oppenlaender, J. 2022. The creativity of text-to-image generation. In *Proceedings of the 25th International Academic Mindtrek Conference*, 192–202.

Paulus, P. B., and Dzindolet, M. 2008. Social influence, creativity and innovation. *Social Influence* 3(4):228–247.

Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.

Sawyer, R. K. 2011. *Explaining creativity: The science of human innovation*. Oxford university press.

Schulman, J.; Zoph, B.; Kim, C.; Hilton, J.; Menick, J.; Weng, J.; Uribe, J.; Fedus, L.; Metz, L.; Pokorny, M.; et al. 2022. Chatgpt: Optimizing language models for dialogue.

Secretan, J.; Beato, N.; D Ambrosio, D. B.; Rodriguez, A.; Campbell, A.; and Stanley, K. O. 2008. Picbreeder: evolving pictures collaboratively online. In *Proceedings of the SIGCHI conference on human factors in computing systems*, 1759–1768.

Shahriar, S. 2022. Gan computers generate arts? a survey on visual arts, music, and literary text generation using generative adversarial network. *Displays* 102237.

Stanley, K. O., and Miikkulainen, R. 2002. Evolving neural networks through augmenting topologies. *Evolutionary computation* 10(2):99–127.

Stanley, K. O.; D'Ambrosio, D. B.; and Gauci, J. 2009. A hypercube-based encoding for evolving large-scale neural networks. *Artificial life* 15(2):185–212.

Stanley, K. O.; Lehman, J.; and Soros, L. 2017. Open-endedness: The last grand challenge you've never heard of. *While open-endedness could be a force for discovering intelligence, it could also be a component of AI itself.*

Stanley, K. O. 2007. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines* 8(2):131–162.

Stanley, K. O. 2019. Why open-endedness matters. *Artificial life* 25(3):232–235.

Tan, W. R.; Chan, C. S.; Aguirre, H. E.; and Tanaka, K. 2017. Artgan: Artwork synthesis with conditional categorical gans. In *2017 IEEE International Conference on Image Processing (ICIP)*, 3760–3764. IEEE.

Tennie, C.; Call, J.; and Tomasello, M. 2009. Ratcheting up the ratchet: on the evolution of cumulative culture. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364(1528):2405–2415.

Tomasello, M.; Kruger, A. C.; and Ratner, H. H. 1993. Cultural learning. *Behavioral and brain sciences* 16(3):495–511.

Tomasello, M. 2009. *The cultural origins of human cognition*. Harvard university press.

Turing, A. M. 1948. Intelligent machinery. Report for National Physical Laboratory. Reprinted in Ince, D. C. (editor). 1992. Mechanical Intelligence: Collected Works of A. M. Turing. Amsterdam: North Holland. Pages 107127. Also reprinted in Meltzer, B. and Michie, D. (editors). 1969. Machine Intelligence 5. Edinburgh: Edinburgh University Press.

Turing, A. M. 1950. Computing machinery and intelligence. *Mind* 59(236):433–460.

Wang, J.; Song, Y.; Leung, T.; Rosenberg, C.; Wang, J.; Philbin, J.; Chen, B.; and Wu, Y. 2014. Learning fine-grained image similarity with deep ranking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1386–1393.