# Tired of choosing? Just add structure and Virtual Reality

Edward Easton(, Ulysses Bernardet, and Aniko Ekart

Department of Computer Science, Aston University, Birmingham, UK {eastonew,u.bernardet,a.ekart}@aston.ac.uk

Abstract. Interactive Evolutionary Computation (IEC) systems often suffer from users only performing a small number of generations, a phenomenon known as user fatigue. This is one of the main hindrances to these systems generating complex and aesthetically pleasing pieces of art. This paper presents two novel approaches to addressing the issue by improving user engagement, firstly through using Virtual Environments and secondly improving the predictability of the generated images using a well-defined structure and giving the user more control. To establish their efficacy, the concepts are applied to a series of prototype systems. Our results show that the approaches are effective to some degree. We propose alterations to further improve their implementation in future systems.

**Keywords:** User Fatigue · Interactive Evolutionary Computation · Virtual Reality.

#### 1 Introduction

Creativity is a part of life that has existed for as long as humankind has existed, from cave paintings to modern-day graffiti artists, it runs deep throughout our very existence [1]. Integral to this is the ability of humans to determine the aesthetic quality of different works of art. Within Evolutionary Art one of the major difficulties is how to effectively formalise these ingrained notions of aesthetics into an automated fitness function. Many approaches have been used to address this problem such as semi-automated systems [2] or fully manual systems [3]. Both types of system try to overcome the problem by replacing, at least partially, the automated fitness function with a human user. These systems are termed Interactive Evolutionary Computation (IEC) systems [3] and this approach has proved effective in a range of different applications such as fashion [4]. Issues exist with both types of system: semi-automated systems still have problems with being able to define user preferences in terms of attributes of the image and fully manual systems are limited by the amount of input received by the user, a limit which is known as user fatigue [3]. The disparity between the number of generations which can be processed between the two types of system is quite high, semi-automated systems can easily complete 60 or more generations [5], whereas the user fatigue limit for fully manual IEC systems is known to be around 10-20 generations. Due to this, user fatigue is a major reason for IEC systems not being able to generate suitably pleasing images. Many novel methods have been introduced in recent years to overcome the limitations enforced by user fatigue such as branching [6], however these methods often enforce other requirements on the system, such as having access to multiple users.

User performance is not a new phenomenon within Computer Science and has often been studied in other application types, often being linked to the concept of user engagement [7]. Within IEC systems, there are two areas where the engagement can be improved.

One of the reasons cited in IEC systems for the presence of user fatigue, is the user is presented with options which do not appear to be related to their previous choices [8]. This can cause the user to stop being engaged with the system and ultimately stop completing generations.

Many IEC systems start with a simplistic, random image and expect the user to provide complexity and a direction to it. This may prevent users from starting to use the application or to quit within the first few generations as the images presented do not engage them sufficiently.

From research conducted in other fields, many different methods of improving user engagement have been suggested such as gamification [9]. One industry which continually encourages users to use their applications for long periods of time and perform a very high number of actions in a single session is the video games industry. This can be seen with users reportedly using the applications for up to 13 hours a week on average [10]. By using techniques to improve the user engagement of the system, we can potentially increase the user fatigue limit. This has already been seen in existing contexts such as the Galactic Arms Race game [11] or Petalz [12] where evolutionary algorithms were applied to content generation which helped to engage users as much as possible.

# 2 Proposed approaches

User engagement is addressed directly within this paper, which introduces two approaches designed to engage users by providing a more predictable experience and presenting artwork to the user in interesting ways. In order to provide a more predictable experience for the user a set structure is described, which the system can follow to present more appropriate images to the user and give them greater control over what is being changed. To give the user new and exciting ways to view the generated artwork, an immersive Virtual Reality environment is presented. These two approaches can be used in isolation and combination to encourage all users to use the system for longer and perform more generations. This section contains details about these approaches. Section 3 looks at how these approaches were realised into a practical application, section 4 contains the experiment design and major results found from evaluating these strategies and finally section 5 presents the conclusions and future work related to these concepts. Full details on the project can be found in [13].

## 2.1 Engagement through Virtual Reality

A major reason to look towards the computer gaming industry when searching for methods to improve user engagement, is that IEC applications share a fair amount of commonality with gaming applications; they both offer a goal for a human user; they both rely on user input to achieve these goals; they both present an interface and control methods for the user to provide input to them (figure 1).



Fig. 1. Commonality between gaming and Interactive Evolutionary Computation (IEC) applications, including the goals, controls and user input aspects which feature in both system types

The gaming industry excels at adopting new technologies and applying them to existing applications. An example of this is the adoption of Virtual Reality (VR) through Head Mounted Displays (HMDs) such as the HTC Vive [14]. Between the market introduction of the first version of the headset in April 2016 and September 2018, over 2500 games have been released for VR on the Steam platform [15]. Though the use of the headset introduces some negative aspects such as its cumbersome nature and its weight, the game industry still offers engaging experiences for VR. This is achieved by utilising the headsets and the unique opportunity VR offers to provide the user with new and exciting ways to experience activities.

There are many aspects which contribute to the user engagement in computer games, such as the storyline and the characters. The application presented in this paper concentrates on a single aspect, the visuals, as a method of improving the user engagement with these systems. This is achieved through the use of VR and presenting the user with images in a 3D context which they can walk around and inspect from multiple different angles. With this added incentive over other evolutionary art applications, we can give the user a reason to want to start and to continue using the system.

#### 2.2 Improving predictability

The lack of predictability within IEC systems can be caused by two aspects, the first is the limited amount of information provided to the system by the user. In most IEC applications the user can only select a small number of preferences and from this information the system cannot determine which aspects of an image the user likes. In addition to this, common data structures such as expression trees have a very low locality [16], where small changes in the representation can lead to large changes in the displayed phenotype. This makes the process of relating elements of the data structure to elements of the produced image very difficult. There are multiple ways which this can be addressed, one method is to use high locality data structures such as the ones presented by Re [17]. However, in doing so the benefits and effects of using other data types are lost.

This paper introduces a set structure to an IEC system in order to improve on the issues mentioned above, whilst still allowing the application to use low locality data structures like expression trees. This gives the user greater control which ultimately leads to the images being amended more predictably for them. The systems are split into 3 different modes, these modes all perform different functions in the process of generating an image. Each one presents multiple images to the user but internally uses different evolutionary processes to change these images over every generation. These three stages are the rough tune mode, the fine tune mode and the lock mode, the relationships between the modes are shown in figure 2.



Fig. 2. State diagram of the transition between modes.

**Rough Tune Mode** The rough tune mode is intended to allow the user to be able to generate a suitable base from which they would like to continue. At this stage the system makes coarse changes to the images without worrying about the outcome. This allows the user to quickly find a base image they like without having to believe it is the best possible image they can create through the system. The user can give a slight direction to the system by selecting their preferred option. Once the user has found an image they think has potential, then they can enter the fine tune mode.

**Fine Tune Mode** The fine tune mode is designed to take the preferred image from the rough tune mode and then make minor adjustments to move it closer towards being the most aesthetically pleasing picture the user believes is possible. Depending on the implementation, this mode can introduce the use of data structures which have a high locality allowing a greater connection between the algorithms used and the user. The mode should be designed so that the user can see the relationship between subsequent sets of images.

Lock Mode The final mode introduced is the lock mode, this mode gives the user a high level of control over what is being changed behind the scenes. It allows the user to select specific regions of a presented image and mark them as not needing any further changes i.e. locked. The system then takes this information and ensures that no further processing will take place on the specified region, until the user manually unlocks it. This mode helps to ensure the system continues to amend images in alignment with the user's intentions, without having to implement complex processes to automatically determine this information. In addition to this, the effective locking of regions within the image effectively reduces the search space the algorithm needs to work through to find the best image for the user.

# **3** Practical Realisation of the Proposed Approaches

To demonstrate the effect these suggestions can have, a series of three systems were created. These systems were comprised of one desktop application and two VR applications, each system being used to represent different levels of immersion. The desktop application represented a no immersion system, the first VR system offered a partially immersive experience being rendered at a low resolution and offering no auditory or haptic feedback to the user. The last VR system was designed to offer a fully immersive experience at high resolution and both auditory and haptic feedback provided. All systems implemented the different modes proposed in this paper, the effectiveness of these can be measured by the comparison of generations completed and the time spent whilst within each system, alongside user feedback. The images were generated using a particle image system similar to ones which has been used in previous applications[18, 19].

### 3.1 Environment Design

**Desktop System** The desktop system was designed with simplicity in mind, where for each mode the user would see a single form displaying the current

images. The system was interacted with using standard desktop controls of context menus and buttons, these allowed the user to select their preferred images, complete generations and lock particles. The layouts of the forms are shown in figure 3.



Fig. 3. Rough Tune and Fine Tune forms in desktop application

The locking of the particles within the desktop environment took the form of allowing the user to cycle through each of the individual particles which combined to construct the image and mark them as being locked or unlocked, as shown in figure 4.



Fig. 4. Desktop locking forms

Virtual Reality Systems To help the user feel a sense of familiarity with the Virtual Environment, it was based around the metaphor of a smart art gallery. The user was placed in a large, open, minimalistic room and the images were presented within frames hanging on the walls. The system used animation to cover the loading times for completing generations, rendering the pictures and to indicate that the gallery itself was amending the images behind the scenes. To allow the system to render the images appropriately each set of items displayed to the user would move around the room in an anti-clockwise direction, where the next set of images could be rendered without interfering with existing ones. Other minor objects were introduced into the environment to help further invoke

the sense the user was in an art gallery such as placing a bench in the middle of the room and having small plaques on the left side of the frames. An overview of the environment is shown in figure 5.



Fig. 5. Art Gallery Environment

In order to present the artwork in new and exciting ways, the VR systems implemented three aspects: the first was allowing the user to move around the environment through teleportation, this let the user easily view the artwork from multiple different angles and from different distances; the system implemented a different version of the lock mode to the desktop environment, where upon entering the mode, the user was presented with a deconstructed view of the particle image, showing each individual particle, as shown in figure 6; finally the system gave the user the ability to take full control over the generated artwork by exporting them as a high quality model into Google's Tilt Brush software[20], an example of a particle image rendered in Tilt Brush is shown in figure 7. Further details about the virtual environments and the Tilt Brush integration can be found at [21]

#### 3.2 Genetic Algorithm

The algorithms in use for the particle systems followed the standard genetic algorithm pattern, using a series of expression trees as the data representation of each image. This algorithm was used in all modes of the systems, however, to make coarser changes to the image within the rough tune mode, only mutation was used. After each generation was completed the entire population was replaced by the children created through the mutation and crossover operations.

Each image was defined as a series of nine expression trees, three were used to determine the initial X, Y, Z position of the particle; three were used to determine the initial RGB colour of the particle and the final three were used to



Fig. 6. Inspecting the individual particles in VR



Fig. 7. Tilt Brush rendered particle model

calculate the X, Y, Z position of the 100 particles over a series of 100 time steps. Due to the integration with Tilt Brush only the X, Y and Z components could be updated and so the entirety of a particle was kept the same colour. Examples of items generated using the particle system are shown in figure 8.

One of three actions could be performed as part of the mutation on a tree, the first action is to entirely replace the node with a new expression but keeping the same child nodes, the second action is to remove a node and all its children entirely and finally just the children of a node could be replaced.

Within the fine tune mode, single point crossover was used by selecting a node at random from both selected parents, the first child was generated by using the tree from the first parent with the selected node replaced by the one from the second parent and the second child is generated by using the tree from



Fig. 8. Example Generated Images



Fig. 9. Example particle image and expression trees

the second parent with the selected node replaced by the one from the first parent. This crossover was applied to each of the 9 trees on each individual.

The use of expression trees as the data structure did cause problems when implementing the locking due to the low locality of the representation. To compensate for this the system took a copy of the 9 expression trees used to generate the locked particle and proceeded to use these trees to plot the path of the particle in all future generations whilst the particle was marked as locked.

Table 1. Parameters used for the rough tune mode

Representation	9 expression trees
Population Size	2
Mutation Operator	0.2 (preferred item), 0.4 (non-preferred item)

Table 2. Parameters used for the fine tune mode

Representation	9 expression trees				
Population Size	4				
Crossover operator	Single point crossover				
Mutation Operator	0.2				

## 4 Experiment and Results

#### 4.1 Experiment design

To investigate the potential effectiveness of the introduced concepts and the created systems, we conducted a small pilot study. Two female and one male participant were given up to 15 minutes to use each of the three systems, being tasked to create the most aesthetically pleasing image they could and to only use the system for as long as they wished to. After the participant had finished using each system, they were asked to complete a short questionnaire, that measured aspects such as their feeling of presence within the system; how effective they believed the system to be and how much they enjoyed using the system. In addition to the questionnaire, the system also recorded user actions, including the mode changes made, the number of generations which were completed, the time spent in each system and how many particles were locked and unlocked.

Within the three systems created, it would be expected that the highest values for user satisfaction and number of generations completed are found on the full immersion VR system and the lowest on the desktop system. This will be a clear indication that all the proposed methods in this report contribute effectively to reducing user fatigue.

## 4.2 Results

The analysis of the usage duration shows that the full fidelity system was being used for almost twice as long as the desktop system (table 3)

The questionnaire data showed that the users rated the full immersion VR system higher, both in terms of how much they enjoyed the interaction, and as to whether the participant would choose to use the system more in future (figure 10).

The ratings showed that all of the systems consistently allowed the participant to feel that they were able to make noticeable progress towards generating an aesthetically pleasing image (figure 11). When asked how excited they were about seeing the next set of images, participants rated the VR higher than the other systems. This indicates that the VR systems did add value to the user and improve engagement.

However, in contrast to these findings, the actual number of generations being completed using the VR system were relatively low, compared to a much higher number completed within the desktop system (table 3).



Fig. 10. User ratings for enjoyment out of 5



Fig. 11. User ratings for ability to make progress, marked out of  $5\,$ 

Table 3.	Basic	usage	statistics	of the	three	systems
----------	-------	-------	------------	--------	-------	---------

	Desktop	Low immersion	Full immersion
Avg. time spent [m:s]	6:03	8:29	10:03
Avg. generations completed $[#]$	35	9.3	8.3

**User feedback** In addition to quantitative results, we received a high amount of positive informal feedback from the participants, especially surrounding the Virtual Reality systems and the lock mode, where the users enjoyed being in the lock mode and viewing the artwork in the deconstructed view.

## 5 Conclusion

The results obtained provide a mixed indication about the systems being tested. Overall each system seemed to work as intended and generally the use of VR and multiple stages within the evolutionary process both improved on the user engagement and predictability of the presented images.

The use of VR seemed to provide three benefits; the user spent longer within the systems, enjoyed using them more and were more interested in seeing the next round of images, all three of these benefits indicate the users were more engaged with the system. However, this did not lead to an increase in the number of generations completed as as shown in table 3. There could be many different reasons for this such as the small size of the pilot study, however in the authors' opinion the implementation of the environments could have been the main cause of this. The virtual environments gave the user a greater set of available activities such as inspecting the images from multiple angles or spending an increased amount of time in the lock mode in order to view the deconstructed image, other aspects which may have reduced the number of generations completed could include the time it took for the animations to show and hide the frames. These aspects of the implementation would all have contributed to severely reducing the number of generations the user could perform in the VR environments and would explain the results we obtained.

The consistency of the rating the users gave all systems for ability to make progress suggests that the inclusion of the three modes in the system positively impacted its ability to provide predictable images to the user. This did not translate into an increase in the number of generations performed, however given the limitations mentioned above, this is perhaps not surprising.

A lot of positive points can be taken from the work presented in this paper, both the use of VR and modes seemed to positively contribute to the improving the engagement of the user. In addition to this, a greater understanding of how to implement these new concepts has been obtained, which allows them to be tested more thoroughly in future to establish their effectiveness and reduce the overhead of the systems to obtain more reliable results.

#### 5.1 Future Work

As mentioned in the previous section, several areas have been identified where the implementation of these ideas could be improved, such as reducing the amount of time a user has to wait to see the next generation. With the information gained from this pilot study, improved implementations of these systems can be created, which can then be used to run a full-scale experiment to get a clearer understanding of the benefits of the introduced concepts.

In addition to further application of these systems to reducing user fatigue in evolutionary art, they offer an exciting base from which non-related research can be completed. Various aspects of these systems cover a wide range of topics within computer science, such as immersion and presence, 3D drawing methods or the investigation of 3D aesthetics.

## References

- Henshilwood, C.S., D'Errico, F., van Niekerk, K.L., Dayet, L., Queffelec, A., Pollarolo, L.: An abstract drawing from the 73,000-year-old levels at Blombos Cave, South Africa. Nature 562(7725), 115–118 (oct 2018), http://www.nature.com/articles/s41586-018-0514-3
- Ekárt, A., Sharma, D., Chalakov, S.: Modelling human preference in evolutionary art. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) (2011)
- 3. Takagi, H.: Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation. Proceedings of the IEEE (2001)
- Loureno, N., Assuno, F., Mas, C., Machado, P.: Evofashion: Customising fashion through evolution. In: International Conference on Evolutionary and Biologically Inspired Music and Art. pp. 176–189. Springer (2017)
- Cohen, M.W., Cherchiglia, L., Costa, R.: Evolving mondrian-style artworks. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) (2017)
- Secretan, J., Beato, N., D Ambrosio, D.B., Rodriguez, A., Campbell, A., Stanley, K.O.: Picbreeder: evolving pictures collaboratively online. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 1759–1768. ACM (2008)
- O'Brien, H.L., Toms, E.G.: What is user engagement? a conceptual framework for defining user engagement with technology. Journal of the American society for Information Science and Technology 59(6), 938–955 (2008)
- Davies, E., Tew, P., Glowacki, D., Smith, J., Mitchell, T.: Evolving atomic aesthetics and dynamics. In: International Conference on Evolutionary and Biologically Inspired Music and Art. pp. 17–30. Springer (2016)
- Sureephong, P., Puritat, K., Chernbumroong, S.: Enhancing user performance and engagement through gamification: Case study of aqua republica. In: Software, Knowledge, Information Management & Applications (SKIMA), 2016 10th International Conference on. pp. 220–224. IEEE (2016)
- 10. Essential Facts About the Computer and Video Game Industry The Entertainment Software Association, http://www.theesa.com/about-esa/essential-factscomputer-video-game-industry/
- 11. Hastings, E.J., Guha, R., Stanley, K.: Evolving content in the galactic arms race video game. In: Computational Intelligence and Games. pp. 241–248. IEEE (2009)
- Risi, S., Lehman, J., D'Ambrosio, D., Hall, R., Stanley, K.: Combining searchbased procedural content generation and social gaming in the petalz video game. In: Aiide. AAAI (2012)
- Easton, E.: Investigating User Fatigue in Evolutionary Art. Masters, Aston University (2018)
- 14. VIVE Buy VIVE Hardware, https://www.vive.com/eu/product/
- 15. SteamSpy All the data and stats about Steam games, https://steamspy.com/
- Galván-López, E., McDermott, J., O'Neill, M., Brabazon, A.: Defining locality as a problem difficulty measure in genetic programming. Genetic Programming and Evolvable Machines 12(4), 365–401 (2011)
- Re, A., Castelli, M., Vanneschi, L.: A comparison between representations for evolving images. In: International Conference on Evolutionary and Biologically Inspired Music and Art. pp. 163–185. Springer (2016)

- Colton, S., Cook, M., Raad, A.: Ludic considerations of tablet-based evo-art. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) (2011)
- 19. Hull, M., Colton, S.: Towards a general framework for program generation in creative domains. Programme Committee and Reviewers p. 137 (2007)
- 20. Tilt Brush by Google, https://www.tiltbrush.com/
- 21. Easton, E.: C# Tilt Brush Toolkit, https://github.com/Prystopia/c-sharp-tiltbrush-toolkit