

Fitness and Novelty in Evolutionary Art

Adriano Vinhas¹, Filipe Assunção¹, João Correia¹,
Aniko Ekárt², and Penousal Machado¹

¹ CISUC, Department of Informatics Engineering,
University of Coimbra, Coimbra, Portugal
{avinhas, fga, jncor, machado}@dei.uc.pt,

² Aston Lab for Intelligent Collectives Engineering (ALICE),
Computer Science, Aston University, Birmingham, United Kingdom
a.ekart@aston.ac.uk

Abstract. In this paper the effects of introducing novelty search in evolutionary art are explored. Our algorithm combines fitness and novelty metrics to frame image evolution as a multi-objective optimisation problem, promoting the creation of images that are both suitable and diverse. The method is illustrated by using two evolutionary art engines for the evolution of figurative objects and context free design grammars. The results demonstrate the ability of the algorithm to obtain a larger set of fit images compared to traditional fitness-based evolution, regardless of the engine used.

Keywords: Novelty Search, Evolutionary Art, Multi-objective Optimisation

1 Introduction

Computational Creativity research posits that an answer to a problem is considered creative if it is both useful and novel³ [1]. Although this definition is not consensual [2], it is the one that better fits the scope of this work. Bearing this in mind, and applying it to images, it was proved that expression-based evolutionary art, introduced by Sims [3], is theoretically able to generate any possible image [4, 5]. In practice, however, the images generated by a given system tend to share the same overall appearance.

Previous works on evolving figurative or ambiguous images have focused on using an expression-based approach [6–9]. Although they achieved interesting results in terms of recognizability of the desired object(s), an important drawback revealed by the analysis of these works is that there are runs where the evolutionary algorithm is not able to evolve images that meet the requirement of resembling the desired object(s). Moreover, there are some runs which are able to generate those images, but the results tend to be very similar to each other, as a consequence of the evolutionary algorithm’s convergence towards an

³ In the context of the present paper, novelty means phenotypic diversity.

optimum. Problem deceptiveness or lack of explicit diversity exploration can be the explanation for these shortcomings.

When evolving Context Free Design Grammars (CFDGs)[10], the grammar representation expressiveness allows the creation of different individuals [11]. As such, CFDGs are capable of generating a family of solutions by combining a small set of production rules. However, during the evolutionary process, after a few generations individuals tend to converge towards an optimum, i.e., most individuals of the final population will share common genetic background and thus will tend to produce similar visual artifacts.

The work described here takes as its starting point these results and analysis, aiming to address the aforementioned shortcomings by proposing an algorithm that takes both novelty and fitness into account. Adopting the traditional definition of creativity, fitness stands for usefulness and phenotypic diversity stands for novelty. The algorithm proposed herein is used to evolve images in two different domains: figurative images and CFDGs. In both cases, when canonical fitness-based evolution is used, the last generations suffer from a lack of diversity. By incorporating novelty-based mechanisms into evolution we expect to increase diversity. The goal of this work is to propose a novelty search algorithm for evolutionary art purposes, and to analyse its capability to generate diverse and fit images, comparing, for each problem, the traditional fitness-based evolution with one that also takes novelty information into account.

The remainder of this paper is structured as follows. Section 2 surveys the state of the art regarding novelty search algorithms; Section 3 describes the proposed algorithm that will be applied to two different problems in Sections 4 and 5. Experimental setup, results and their analyses are focused upon in Section 6. In Section 7 conclusions are drawn and future work is addressed.

2 State of the Art

Novelty search has become a trending topic in optimisation problems, where it has been used in fields such as robotics [12–14] and arts [15, 16]. Even before the first formalisation of a novelty search algorithm, by Lehman and Stanley [17], there were already reports in the literature concerning the evolution of images using a novelty metric [18]. In the work of Saunders et al. [18], the novelty concept is mathematically modeled according to an interestingness definition based on two factors: (i) the ability to create artifacts out of the box (unexpectedness) and (ii) the feasibility of taking an action as a consequence of the discovery (actionability) [19]. This metric is then used in the context of an Evolutionary Algorithm (EA) designed to evolve novel imagery artworks. In essence, these abstract images were generated using an EA, in which the evaluation was performed by an agent in two steps. First, a self organizing map (SOM) [20] is used to determine the category where an image best fits. The novelty is computed as the classification error of an image being associated to the best category. The second step consists in applying the interestingness function, computing the interest of each image, given its novelty degree.

As mentioned before, the more recent concepts of novelty search follow the steps of Lehman and Stanley’s work, which propose a novelty search algorithm that aims at valuing each image’s uniqueness. For each individual, a novelty score is computed, taking into account its neighbours and an archive containing the most novel individuals. Each novelty score computation requires a phenotype comparison, using a dissimilarity metric, between the individual being evaluated and a set of neighbours chosen from the population and the archive. Then, the novelty score for the individual being evaluated (ind_{eval}) is defined as the average of the dissimilarity scores of the k most similar neighbours, as in:

$$\text{nov}(\text{ind}_{eval}) = \frac{1}{k} \sum_{j=1}^k \text{dissim}(\text{ind}_{eval}, \text{ind}_j), \quad (1)$$

where *dissim* denotes the chosen dissimilarity metric, and j the j -th most similar individual (ind_j) when compared to ind_{eval} .

In order to enable the use of novelty search techniques for constrained problems, Stanley and Lehman proposed the Minimal Criteria Novelty Search (MCNS) algorithm [21]. The idea of employing novelty search *per se* is not suitable for this kind of constrained problems because the function used to evolve individuals enhances diversity without any boundaries and hence, creates unsuitable solutions. Therefore, MCNS tries to tackle this problem by decimating the population, i.e., any valid solution would be evaluated using the respective novelty function and unsuitable solutions would be assigned zero fitness.

Liapis et. al based their work on MCNS [22]. Although MCNS does not favour the creation of unsuitable solutions, the idea of decimating a population means that one is assigning the same value (zero) to all unsuitable solutions. As the initial generations of EAs usually hold lots of unsuitable solutions, the selection process would become similar to a random search. This drawback becomes more evident if the problem that one is trying to solve is highly constrained. In order to tackle this disadvantage, Liapis et. al proposed two solutions based on the feasible-infeasible two-population genetic algorithm (FI-2pop GA) [23]. In essence, the idea of FI-2pop GA is to evolve valid solutions and unsuitable solutions as if they are two different populations. However, there can be exchanges of individuals between populations, because a valid solution can become unsuitable and vice-versa. The two solutions proposed were: (i) a feasible-infeasible novelty search (FINS), which consists in evolving the feasible population with novelty search as in the work of Lehman and Stanley[17], while evolving the infeasible one with the goal of minimising each individual’s distance to the feasibility border, and (ii) a feasible-infeasible dual novelty search (FI2NS), which consists in applying novelty search as in [17], in both populations separately, using two different novelty archives. Their results suggest that both solutions are superior to the MCNS performance but, while FINS is able to get feasible individuals within a lower number of generations, FI2NS creates a more diverse set of feasible solutions.

Despite the good performance of novelty approaches, authors tend to point out their lack of ability to maintain good performance in problems with big

search spaces. The rationale is that novelty approaches saturate the search space, until they eventually get a reasonable or optimal solution. However, in big search spaces this saturation is much harder to achieve, and the solutions will fail to reach an optimal solution or, in an optimistic scenario, the desired solution is found after a large number of generations. This problem motivated the development of hybrid mechanisms, which use both fitness and novelty-based functions to guide evolution. Within this scope, we highlight three works: [12, 24, 14].

In Krcah et al. [12], novelty search is applied during a predefined number of generations, and the remainder of the process is performed with a goal-based evolution. Although it is a simple solution, it will not avoid a later convergence to a local optima. In Cuccu et al. [24], novelty and fitness are merged using a weighted combination approach. This solution requires prior knowledge about the weight values, which depend on the domain. Finally, Mouret et al. [14] used a Pareto-based multi-objective evolutionary algorithm to combine fitness and novelty, in order to help a robot which is subject to a maze navigation task. Even though it is an interesting idea, preliminary tests indicate that it is not appropriate in an evolutionary art context since finding phenotypically diverse images is significantly easier than finding fit images and the search space is vast. Thus, in this conditions evolution will produce a seemingly endless stream of images that are “novel” but not fit.

3 Proposed Algorithm

In this section, a new novelty search algorithm is proposed, designed to evolve a diverse set of suitable images. Therefore, the main goal of this algorithm is to generate a more phenotypically diverse set of images than the set that would be created by a traditional fitness based EA. In essence, it is a method capable of evolving images according to two criteria that are chosen automatically by analysing the quality of the images produced in each generation. One criterion is to look for the best images according to a fitness function and the other consists in taking novelty and fitness as two different objectives to be maximised. As previously mentioned, considering fitness or novelty alone is not suitable for the problem at hand [25].

The algorithm’s flowchart is similar to the traditional EA one, differing only in two main aspects: (i) the creation of an archive to store the most novel solutions and, (ii) a customised selection mechanism which is able to consider single or multiple objectives using a tournament based strategy. The algorithm’s flow is shown in Figure 1, and can be summarised as follows:

1. Randomly initialise the population;
2. Render the images (phenotypes) from the individuals’ genotypes;
3. Apply the fitness function to the individuals;
4. Select the individuals that meet the criteria to be in the archive (archive assessment);
5. Select the individuals to be used in the breeding process. The individuals are picked using one of the following criteria: (i) according to their fitness,

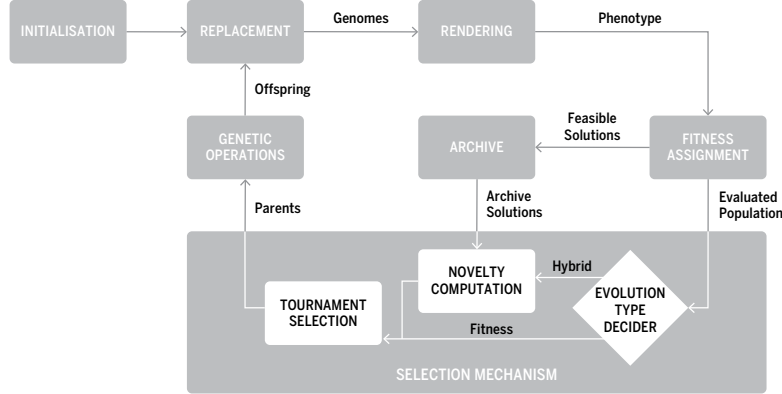


Fig. 1: Flow of the hybrid algorithm proposed.

- as a standard EA; (ii) taking into account both the fitness and the novelty metric, computed using the archive members;
6. Employ genetic operators to create the new generation of solutions, that will replace the old one;
 7. Repeat the process starting from step 2, until a stop criterion is met.

3.1 Archive Assessment

In this work, the archive has an unlimited size and it plays an important role, because it is used to evaluate our solution and prevents the algorithm from exploring areas of the search space already seen before. The idea is that the archive should represent the spectrum of images found to date, and for this reason, the bigger the archive is, the more capable is the algorithm of generating suitable and diverse images. Whereas in the previously mentioned works the archive size is limited, we opted for not restricting it.

At this stage, a candidate individual has its fitness assigned and it has to meet two requirements in order to be added to the archive: (i) its fitness must be greater or equal than a threshold f_{min} , defined by the user; (ii) it needs to be different from those that already belong to the archive. This process is performed by computing the average dissimilarity between the candidate and a set of k -nearest neighbours. When the average dissimilarity is above a predefined dissimilarity threshold, $dissim_{min}$, the individual is added to the archive. The dissimilarity metric for an image i is computed as:

$$dissim(i) = \frac{1}{\max_{arch}} \sum_{j=1}^{\max_{arch}} d(i, j), \quad (2)$$

where \max_{arch} is a predefined parameter which represents the number of most similar images to consider when comparing with image i , and $d(i, j)$ is a distance

metric that measures how different two images (i and j) are. From this dissimilarity measure there are two exceptions that should be highlighted. If there are no entries in the archive, the first individual that has a fitness above f_{min} is added. Moreover, if the number of archive entries is below max_{arch} , Equation 2 is used with the number of archive entries instead of max_{arch} .

3.2 Selection Mechanism

The selection mechanism is important to shape how evolution will proceed, depending on the results obtained in a given generation. Our novelty approach has a customised selection mechanism which can switch between a fitness-based strategy and a hybrid mechanism that considers both fitness and novelty. It starts as a fitness guided evolution; however, that can change according to a decision rule, which is described as:

$$\begin{cases} \text{change_to_fitness,} & \text{feasible}_{inds} < T_{min} \\ \text{change_to_hybrid,} & \text{feasible}_{inds} > T_{max}, \end{cases}$$

where $feasible_{inds}$ is the number of individuals of the current generation that have a fitness above the threshold f_{min} , T_{min} is the threshold used to verify if evolution should be changed to fitness, and T_{max} is used to verify if it should be changed to hybrid.

In fitness guided evolution, the tournament selection is based on the fitness values of the candidate solutions, as in a standard EA. If hybrid evolution is chosen, it is necessary to compute the novelty of each selected individual, and perform a Pareto-based tournament selection, using the novelty and fitness of each selected individual as two different objectives to maximise.

The novelty computation process is inspired by Lehman and Stanley’s work, as described in Equation 1, with one small change: the k most similar images are considered from the set of the selected individuals and the archive, instead of considering the whole population and the archive. An example of this novelty computation is illustrated in Figure 2. In Figure 2, considering $k = 4$ and a tournament size of 5, the dashed lines denote the chosen individuals to compute novelty, and it is possible to see that from the 4 nearest individuals picked, 3 were chosen from the tournament while the remaining one was chosen from the archive.

At this stage, each selected individual has a fitness and novelty value, and there is the need to determine the winner of the tournament. This process is inspired by multi-objective EAs, namely the Pareto-based approaches, which select the best individuals based on their dominance or non-dominance when compared to other individuals [26]. In this work, the hybrid tournament selection determines the non-dominant solutions by comparing, among the selected individuals, on the basis of both fitness and novelty. After computing the set of non-dominant individuals, we have the so-called Pareto front. The tournament winner will be selected by randomly retrieving one of the solutions of the Pareto front.

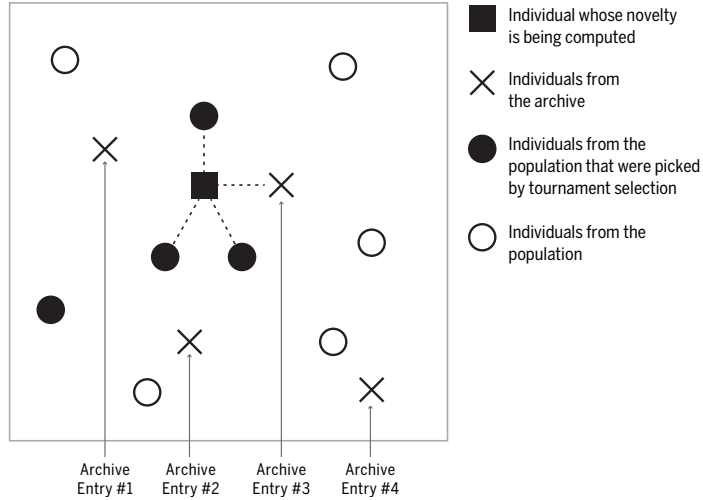


Fig. 2: Novelty computation for an individual.

4 Evolving Figurative Images

One of the problems for which we decided to apply the proposed algorithm was the evolution of figurative images. In this problem, an expression-based approach is used to build the individuals' genotypes and the evolution is guided with the help of an external object classifier. Previous works tackled the evolution of singular objects such as faces, lips, leaves or flowers [6–8], and the evolution of several objects at the same time, attempting to create ambiguous images [9].

For this work, we used the **geNeral purpOse expREssion BAsed EVolutionary aRt Tool** (norBERT), which is able to evolve figurative images using several evolution strategies [25]. It uses a tree representation to encode individuals and create images from those trees, using a rendering process which consists in generating an output value for each image pixel, which represents its colour.

For the fitness assignment task, an object detector is used to assess the images' quality. This object detector is a cascade classifier based on the work of Viola and Jones [27]. The architecture of this cascade classifier is shown in Figure 3. The cascade classifier is composed of several stages and each stage tests the presence of a group of low level features (Local Binary Pattern features [28]).

The fitness function is built by obtaining internal information from the cascade classifier and combining them as follows:

$$f(x) = \sum_i^{\text{cstages}_x} (\text{stagedif}_x(i) \times i) + \text{cstages}_x \times 10, \quad (3)$$

where cstages_x is the number of stages that an image x passed and stagedif_x is the difference between the score obtained in a given stage by an image x and the

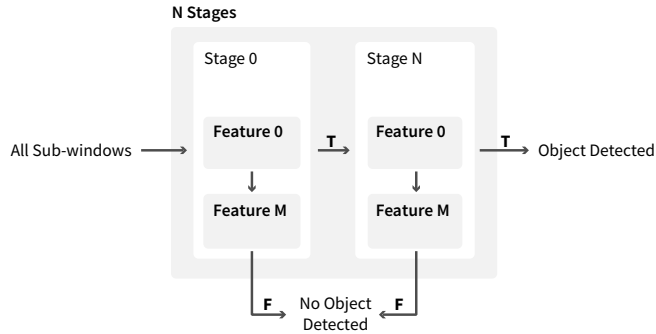


Fig. 3: Cascade classifier structure used for object detection.

threshold necessary to pass that stage. The detailed description of the problem of evolving figurative images, and specifically the object detection method and the classifier training process can be found in [6].

5 Evolving Context Free Design Grammars

In the last section we studied how novelty search mechanisms can be applied for the evolution of figurative images. Aiming at better understanding the impact of novelty mechanisms applied to a broader scope of evolutionary algorithms, we applied this kind of procedure to a different environment, which works as a proof of concept for the ability of the methodology to properly evolve solutions under different domains.

Taking into account the previous statements, we address the evolution of CFDGs [10], which, in simple words, are a powerful way of generating images through a compact set of rules (for full description check [10]). CFDGs are mapped into images by using a software tool named Context Free Art [29].

The evolutionary engine used for this task is an extension of the one discussed in [30] and [31], which has been generalised, in order to enable it to deal with every domain that can be represented by means of a formal grammar (in the Backus-Naur Form). It is thoroughly described in [32] and the application to the domain of CFDGs, along with a large set of comprehensive experiments, is detailed in [11].

The purpose of the engine is to evolve individuals expressed from grammar formulations. In essence, a grammar is formed by a 4-tuple: (V, Σ, R, S) where: V is a set of non-terminal symbols; Σ is a set of terminal symbols; R is a set of production rules that map from V to $(V \cup \Sigma)^*$; S is the initial symbol. Additionally, grammars can also be considered augmented, because we allow the specification of parameters in the calls to terminal and non-terminal symbols. One can also say that the grammars are non-deterministic, because it is possible to define the same non-terminal symbol more than once. When several production rules are applicable, one of them is randomly selected and the expansion

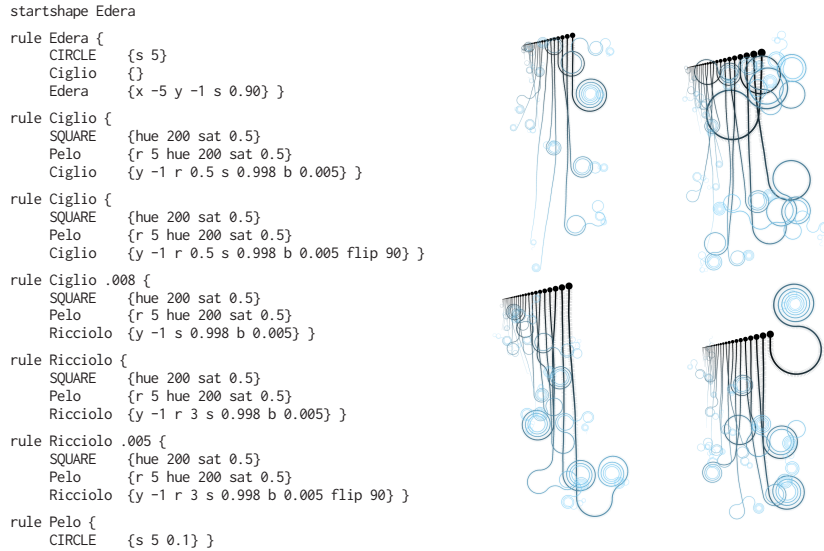


Fig. 4: CFDG adapted from www.contextfreeart.org/gallery/view.php?id=165 together with examples of images produced by its rendering.

proceeds. Figure 4 depicts an example of a CFDG and several images that can result thanks to the probabilistic nature of its associated rendering algorithm. More information regarding representation and genetic operators for CFDGs can be found in [11, 32].

To test the influence of novelty in evolution, we adapted the methodology described in Section 3 to this evolutionary engine. To assess the quality of each individual we use a combination of aesthetic measures: contrasting colours [11] and Bell Curve fit. Bell Curve fit [33] is based on the observation that fine-art works tend to exhibit a normal distribution of colour gradients. Thus, we calculate the deviation from the Gaussian normal distribution of colour gradients of the individuals. Contrasting colours counts the number of colours present in an image, discarding similar ones. Therefore, the engine aims at promoting the evolution of images that depict a wide range of colours, possessing a normal distribution of their gradients.

6 Results

In the current section, we describe the set of experiments that were conducted using the two different engines summarised in Sections 4 and 5, respectively for the evolution of figurative images and CFDGs. In Section 6.1 we present the setup used for each problem, while Section 6.2 details the results. For each problem, we compare the performance of a standard fitness evolution approach and our proposed algorithm, which takes novelty into account.

Table 1: Exerimental parameters.

Parameter	Figurative Images	CFDGs
Number of runs	60	30
Number of generations	500	500
Population size	100	100
Image size	64×64	128×128
f_{min}	Flower detected	0.52
$dissim_{min}$	20%	20%
Dissimilarity metric	RMSE ⁴	RMSE
(T_{min}, T_{max})	(2, 15)	(2, 15)
max_{arch}	5	5

6.1 Experimental Setup

Table 1 specifies the parameters used to perform the experiments for evolving figurative art images and CFDGs. We performed 30 runs for the problem of evolving CFDGs and 60 runs in the evolution of figurative images.

The evolution of figurative images sometimes results in an empty archive. Therefore, a larger number of runs were needed to ensure a sufficient number of runs with non-empty archive. The entry condition for the archive is different in the two cases. In norBERT, classifiers are used to detect objects and therefore, a feasible solution would be one that contains an object. In the current section, we conducted tests aiming to get images that resemble flowers. In the case where CFDGs are evolved, a fitness threshold is defined for feasible solutions (in the current scenario, 0.52, which was empirically determined).

The remaining parameters, i.e., those specific to the evolution of figurative images and CFDGs can be found, respectively, in [25] and [32].

6.2 Results

The experiments serve as basis for comparing fitness-based and hybrid guided evolutions for figurative images and CFDGs. For both domains, we use the average fitness of the best individuals (Figure 5), archive sizes (Figure 6) and populational RMSE (Figure 7) as comparison metrics. The evolution towards feasible zones of the search space – i.e., containing images recognized as flowers when evolving figurative images, or with a fitness above the considered threshold when evolving CFDGs – was observed in both problems. However, there is clearly a difference between the two problems in terms of difficulty. While in the evolution of CFDGs the f_{min} constant is determined empirically and based on subjective options, in the evolution of figurative images the constraint is more difficult to match, as images need to resemble an object. Figures 6 and 7 highlight this situation, as the archive sizes are higher in the CFDGs case and populations are more diverse. This suggests that, when evolving CFDGs, feasible zones of the search space are more accessible.

⁴ RMSE stands for Root Mean Squared Error.




Fig. 5: Evolution of fitness across generations. On the left for figurative images (averages of 60 independent runs), on the right for CFDGs (averages of 30 independent runs).

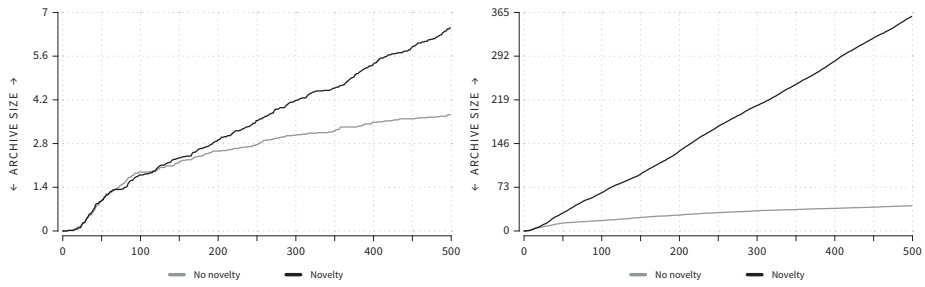


Fig. 6: Evolution of the archive size across generations. On the left for figurative images (averages of 60 independent runs), on the right for CFDGs (averages of 30 independent runs). Be aware of the difference in scale.




Fig. 7: Evolution of population dissimilarity across generations. On the left for figurative images (averages of 60 independent runs), on the right for CFDGs (averages of 30 independent runs).

Regardless of the difficulty of the problem, we aim to demonstrate that the proposed hybrid algorithm is able to outperform the fitness-only approach in terms of generating diverse images, with some cost associated in the evolutionary process, due to additional similarity calculations and possibly more generations needed for convergence.

Concerning the behaviour of fitness, we analysed how best fitness evolves over the generations (see Figure 5). Results are normalised, by dividing by the best fitness value found throughout all evolutionary runs. The differences in the evolution of figurative images are too small to be seen. In the CFDGs case, the hybrid algorithm has more difficulty to converge towards its best solution, as the difference in convergence between generation 25 and 250 shows. After generation 250, fitness values from both hybrid and fitness evolutions become very close. This is due to the fact that evolving using a novelty metric blurs the search for the fittest individuals. The evolution process can thus move towards new areas, allowing us to get more unique individuals, at the cost of lower fitness values. Our main goal is not to maximize fitness *per se*, instead, we are interested in obtaining good individuals (the ones that are above f_{min}) that are as different from each other as possible. The impact of using a novelty metric in our hybrid algorithm in terms of number of good and unique individuals (those that enter the archive) varies according to the problem difficulty. Figure 6 makes the previous statement clear, by highlighting the huge difference between the archive sizes of the two problems, also indicating this way that the evolution of figurative images is harder than CFDGs.

In both cases, the use of novelty metrics in the hybrid algorithm, represented by the darker line, clearly favours the growth of the archive. The difference between fitness-based and hybrid evolution is larger in the case of the simpler problem, where it is possible to get, on average, roughly 360 additional diverse images per run when using the hybrid algorithm. In the case of evolving figurative images, the difference is noticeable but smaller (approximately 3 images).

We have established that novelty search results in larger archives, but we do not know how the algorithm proceeds in terms of phenotypic diversity in a broader scope. In order to analyse how similar the population is over generations, we used the RMSE metric to measure diversity at a population level, by computing the difference between every pair of individuals. As the results depicted in Figure 7 show, for both evolutionary engines, the population diversity using novelty is higher than without it. Since the data does not follow a normal distribution, we use the non-parametric Wilcoxon signed-rank test. All the tests were performed with a confidence level of 95%, and confirm that all the differences reported herein are statistically significant. Moreover, almost across all the generations, the difference between the population RMSE, with and without novelty is higher when evolving CFDGs. This can also be attributed to the fact that it is easier to evolve CFDGs than figurative images.

When evolving figurative images, diversity with novelty converges faster (approximately to 0.14) than when evolving CFDGs, which seems to increase till approximately 0.31. This is expected due to the nature of the problems and to the parameters that were chosen. While evolving figurative images, tournament changes from fitness only to hybrid are much more frequent than when evolving CFDGs. That also explains why at the beginning the diversity of CFDGs decreases. Initially, it is necessary to generate the minimum number of feasible solutions (CFDGs superior to a defined threshold of fitness), sacrificing diver-




Fig. 8: Archive of a run evolving figurative images guided only by fitness.




Fig. 9: Archive (sampled) of a run evolving figurative images with the hybrid method.




Fig. 10: Archive of a run evolving CFDGs guided only by fitness.




Fig. 11: Archive (sampled) of a run evolving CFDGs with the hybrid method.

sity. Until the archive has at least one entry, the two methods are practically identical. Then, after changing to hybrid evolution, the novelty metrics promote

an increase in population diversity, which is also noticeable by the growth of the archive, depicted in Figure 6.

We consider it essential to analyse the results from two different perspectives: the computational one and the human one. While the former is covered by the RMSE-based analysis, the latter was explored through a subjective human analysis. Overall, we can conclude that not only the number of resulting images in the archive is larger in the cases when we use the hybrid algorithm, but also that the differences between them are more visually noticeable. This is confirmed by the archive images of a single representative run starting from the same initial population of images, using both approaches, for the problem of evolving figurative images (Figures 8 and 9) and for evolving CFDGs (Figures 10 and 11).

7 Conclusions and Future Work

In this paper we proposed a new novelty search approach within the context of evolutionary art, and applied it to the evolution of figurative images and CFDGs, using two different evolutionary engines. Our approach is an hybrid algorithm with adapted evolution, in the sense that it changes between fitness based and hybrid evolution, taking into account fitness and novelty as two different objectives to maximise. Tests were designed to compare the performance of our novelty search approach with canonical fitness based evolution.

The experimental results show that our algorithm is able to promote the discovery of a wide set of phenotypically diverse and fit solutions to the problems considered, outperforming fitness based evolution in terms of the number and the diversity of the generated solutions.

Next steps to this work include the study of the conditions used to change the evolution method. At this moment, the method depends on two user-defined parameters. Choosing the best parameter settings requires some insight from the user, which may not be easy to obtain. Domain knowledge may be needed to understand what are the best parameter settings for each problem. We believe that a self-adaptive mechanism that automatically detects when to change the evolution method could be beneficial to this work. Additionally, the application of these techniques to other types of problems, outside the scope of evolutionary art, is also under way.

Acknowledgments

The project ConCreTe acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant number 611733. This research is also partially funded by: Fundação para a Ciência e Tecnologia (FCT), Portugal, under the grant SFRH/BD/90968/2012. The authors also acknowledge the feedback and provided by the blind reviewers of this paper.

References

1. Boden, M.A.: The creative mind: Myths and mechanisms. Psychology Press (2004)
2. Kowaliw, T., Dorin, A., McCormack, J.: Promoting creative design in interactive evolutionary computation. *IEEE transactions on evolutionary computation* **16**(4) (2012) 523
3. Sims, K.: Artificial evolution for computer graphics. *ACM Computer Graphics* **25** (1991) 319–328
4. Machado, P., Cardoso, A.: All the truth about NEvAr. *Applied Intelligence, Special Issue on Creative Systems* **16**(2) (2002) 101–119
5. McCormack, J.: Facing the future: Evolutionary possibilities for human-machine creativity. In: *The Art of Artificial Evolution*. Springer Berlin Heidelberg (2008) 417–451
6. Correia, J., Machado, P., Romero, J., Carballal, A.: Evolving figurative images using expression-based evolutionary art. In: *Proceedings of the fourth International Conference on Computational Creativity (ICCC)*. (2013) 24–31
7. Machado, P., Correia, J., Romero, J.: Expression-based evolution of faces. In: *Evolutionary and Biologically Inspired Music, Sound, Art and Design – First International Conference, EvoMUSART 2012, Málaga, Spain, April 11-13, 2012*. Proceedings. Volume 7247 of *Lecture Notes in Computer Science.*, Springer (2012) 187–198
8. Machado, P., Correia, J., Romero, J.: Improving face detection. In Moraglio, A., Silva, S., Krawiec, K., Machado, P., Cotta, C., eds.: *Genetic Programming - 15th European Conference, EuroGP 2012, Malaga, Spain, April 11-13, 2012*. Proceedings. Volume 7244 of *Lecture Notes in Computer Science.*, Springer (2012) 73–84
9. Machado, P., Vinhas, A., Correia, J.a., Ekárt, A.: Evolving ambiguous images. In: *Proceedings of the 24th International Conference on Artificial Intelligence. IJCAI'15, AAAI Press* (2015) 2473–2479
10. Horigan, J., Lentzner, M.: Context Free Design Grammar version 2 syntax. http://www.contextfreeart.org/mediawiki/index.php/Version_2_Syntax (2015)
11. Machado, P., Correia, J., Assunção, F.: Graph-Based Evolutionary Art. In Gandomi, A., Alavi, A.H., Ryan, C., eds.: *Handbook of Genetic Programming Applications*. Springer, Berlin (2015)
12. Krcak, P., Toropila, D.: Combination of novelty search and fitness-based search applied to robot body-brain co-evolution. In: *Czech-Japan Seminar on Data Analysis and Decision Making in Service Science*. (2010) 1–6
13. Methenitis, G., Hennes, D., Izzo, D., Visser, A.: Novelty search for soft robotic space exploration. In: *Proceedings of the 2015 on Genetic and Evolutionary Computation Conference. GECCO '15, New York, NY, USA, ACM* (2015) 193–200
14. Mouret, J.B.: Novelty-based multiobjectivization. In: *New Horizons in Evolutionary Robotics*, Springer Berlin/Heidelberg (2011) 139–154
15. Liapis, A., Yannakakis, G.N., Togelius, J.: Sentient sketchbook: Computer-aided game level authoring. In: *FDG*. (2013) 213–220
16. Secretan, J., Beato, N., D'Ambrosio, D.B., Rodriguez, A., Campbell, A., Folsom-Kovarik, J.T., Stanley, K.O.: Picbreeder: A case study in collaborative evolutionary exploration of design space. *Evolutionary Computation* **19**(3) (2011) 373–403
17. Lehman, J., Stanley, K.O.: Exploiting open-endedness to solve problems through the search for novelty. In: *Proc. of the Eleventh Intl. Conf. on Artificial Life (ALIFE XI)*, Cambridge, MA, MIT Press (2008)

18. Saunders, R., Gero, J.S.: The digital clockwork muse: A computational model of aesthetic evolution. *Proceedings of the AISB* **1** (2001) 12–21
19. Silberschatz, A., Tuzhilin, A.: What makes patterns interesting in knowledge discovery systems. *Knowledge and Data Engineering, IEEE Transactions on* **8**(6) (1996) 970–974
20. Kohonen, T.: *Self-organization and Associative Memory: 3rd Edition*. Springer-Verlag New York, Inc., New York, NY, USA (1989)
21. Lehman, J., Stanley, K.O.: Revising the evolutionary computation abstraction: minimal criteria novelty search. In: *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, ACM (2010) 103–110
22. Liapis, A., Yannakakis, G., Togelius, J.: Enhancements to constrained novelty search: Two-population novelty search for generating game content. In: *Proceedings of Genetic and Evolutionary Computation Conference*. (2013)
23. Kimbrough, S.O., Koehler, G.J., Lu, M., Wood, D.H.: On a feasible-infeasible two-population (FI-2Pop) genetic algorithm for constrained optimization: Distance tracing and no free lunch. *European Journal of Operational Research* **190**(2) (2008) 310 – 327
24. Cuccu, G., Gomez, F.: When novelty is not enough. In: *Proceedings of the 2011 International Conference on Applications of Evolutionary Computation - Volume Part I. EvoApplications'11*, Berlin, Heidelberg, Springer-Verlag (2011) 234–243
25. Vinhas, A.: *Novelty and figurative expression-based evolutionary art*. Master's thesis, Department of Informatic Engineering, Faculty of Sciences and Technology, University of Coimbra (July 2015)
26. Fonseca, C.M., Fleming, P.J.: An overview of evolutionary algorithms in multiobjective optimization. *Evolutionary computation* **3**(1) (1995) 1–16
27. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on* **1** (2001) 511
28. Ojala, T., Pietikäinen, M., Harwood, D.: A comparative study of texture measures with classification based on feature distributions. *Pattern Recognition* **29**(1) (1996) 51–59
29. Horigan, J., Lentczner, M.: *Context Free*. <http://www.contextfreeart.org/> (2014)
30. Machado, P., Nunes, H.: A step towards the evolution of visual languages. In: *First International Conference on Computational Creativity*, Lisbon, Portugal (2010)
31. Machado, P., Nunes, H., Romero, J.: Graph-based evolution of visual languages. In Chio, C.D., Brabazon, A., Caro, G.A.D., Ebner, M., Farooq, M., Fink, A., Grahl, J., Greenfield, G., Machado, P., O'Neill, M., Tarantino, E., Urquhart, N., eds.: *Applications of Evolutionary Computation, EvoApplications 2010: EvoCOMNET, EvoENVIRONMENT, EvoFIN, EvoMUSART, and EvoTRANSLOG*, Istanbul, Turkey, April 7-9, 2010, Proceedings, Part II. Volume 6025 of *Lecture Notes in Computer Science.*, Springer (2010) 271–280
32. Assunção, F.: *Grammar based evolutionary design*. Master's thesis, Department of Informatic Engineering, Faculty of Sciences and Technology, University of Coimbra (July 2015)
33. Ross, B.J., Ralph, W., Zong, H.: Evolutionary image synthesis using a model of aesthetics. In: *Evolutionary Computation, 2006. CEC 2006. IEEE Congress on, IEEE* (2006) 1087–1094