

# Serendipitous evolution; In pursuit of neglected beauty in evolutionary algorithms.

Marian Maas<sup>a</sup> Menno van Elk<sup>b</sup> Tiago Borges Coelho<sup>a</sup>

<sup>a</sup> *mpf.maas@gmail.com*

<sup>b</sup> *Menno.vanElk@phil.uu.nl*

<sup>c</sup> *borges@consola.org*

## Abstract

Current evolutionary methods in computer science employ fitness functions to impel the specialization of virtual organisms. Throughout the generations, the variants that are not near a local optimum get discarded, to allow for more able mutations to succeed. Our research purpose is to analyze these least fit variations, in order to produce new perspectives and methodologies for evolutionary computing. Our approach is to run evolutionary algorithms in visual simulations and report on serendipitous behaviors, while making conjectures about why they occurred.

## 1 Introduction

The concept of evolution is but a learning machine engineered to climb hills (Minsky, 1994) [1]. These hills arise from fitness functions that shape entire populations towards a goal. But in nature, goals are never a conscious factor for evolution, we are never conscious of what we are turning into. We evolve to no particular goal, even as we try hard to adapt the environment to us, in the elusive attempt to stop evolution at some optimum. The purposeless nature of evolution is a feature that should be taken into serious consideration, as it allows for variation, which in turn assures genetic quality. The same holds for specialization, as noted on the sentence “the specialist is he who makes no small mistakes while moving towards the Grand Fallacy” (McLuhan, 1967). As the fitness function specializes the purpose of a population, it produces behaviours that are considered less fit, but are pragmatic proof of extraordinary accomplishments by evolution. Our interest lies on understanding what will become of a population of least fit individuals, if they are allowed to persist. As generations pass, they should evolve towards some recognizable behaviour, be it closer or further away from the goal imposed by the fitness function. Alternatively we speculate the viability of goal-less environments to produce results that can be compared to the outcome of fitness-based evolutionary methodologies.

This research was inspired by what Minsky coined as Negative Expertise: the importance of knowing what not to do to be successful in any particular task. In his article [1] about negative expertise Minsky stressed how species are at a local optimum in their current environment. According to biologists most mutations are deleterious; no additions are made to the physiology of organisms anymore, there is just deletion of existing mutations. Our position at the top of the evolutionary hill poses an interesting question: What is further down the hill?

The less fit species must have had some rather bad ideas over time, to make them produce little or no offspring. The fact that those ideas were bad is proven by the fact

that the organisms didn't turn out to reproduce much. This is not necessarily true though: evolution is ruled by probability for the major part; fitter individuals have a better chance to survive, but there's no guarantee that only the wonderfully good ideas stay in existence. There might be something good about evolutionary unsuccessful species.

In our research we view the process of evolution as a conventional censor and see if there's beauty in the alternatives that are discarded because of lack of fitness. Because as Minsky explains: to be a creative artist means that you can (or even must) generate some unconventional alternatives to what the audience thinks is beautiful, in order to unsuppress their conventional censors, because "beauty is neither in the eye, nor even the mind of the observer, but precisely the opposite: it may lie in the power to inactivate many of that observer's internal critics." (Minsky, 1994)

Exactly what the result of our research should be we do not know, we just want to look into the phenomena of neglected beauty and discarded goodness.

## 2 Description

For our experiments we chose genetic algorithms to simulate evolution in a virtual environment. A genetic algorithm is a search technique used in computing to find true or approximate solutions to optimization and search problems. They are implemented as a computer simulation in which a population of abstract representations of candidate solutions (called individuals, creatures) to an optimization problem evolves toward better solutions. Its mechanisms are inspired by biological evolution: reproduction, mutation, recombination, natural selection and survival of the fittest.

Genetic algorithms are not the most efficient problem solving mechanisms around though. Russel and Norvig [2] even doubted whether the preference of some programmers to experiment with them was purely on rational grounds: "At present, it is not clear whether the appeal of genetic algorithms arises from their performance or from their aesthetically pleasing origins in the theory of evolution." (2003)

Especially this origin in the theory of evolution is what makes them useful for our experiments though, for it's exactly this process that we are interested in. We do not want to test the functionality of genetic algorithms, but we want to simulate evolution to study what happens when conventional variables are adapted.

The one we used for our first experiment is called "Genetic Algorithm; Evolving Shakespeare" and was initially made by Daniel Shiffman [4]. It started out with a 100 random strings of signs for individuals, with the goal to create the sentence 'to be or not to be'. The fitness of the individuals was measured by how close they were to the goal, every right sign increasing their value. This fitness function was then used to calculate the chances of reproduction for the individual; the higher the fitness the bigger the chance to end up in the mating pool. The next generation would consist of 100 strings again, some of them similar to their parents and some with mutations, all with their own fitness. This would go on for a number of generations and eventually quit when one of the individuals reached the goal.




Fig. 1 The added graphical interface to Shiffman's Genetic Algorithm; Evolving Shakespeare.

We adapted this program to see what happens when the fitness function has less influence over reproduction, we also added a graphical representation (fig. 1) of all individuals in the current generation, instead of just showing the current most successful string.

The second experiment was done using the “Disney meets Darwin” algorithm made by Jeffrey Ventrella [5]. This application joins the art of character animation, or the endowment of expressive movement to figurative illustration, with the science of genetic evolution. As there are rewarding functions for features such as movement speed and head bounciness, the creatures evolve towards desired goals (speed, bounciness). What made us choose this environment is the fact that it can be customized, so the user can effectively control the genetic code of the individuals. Also of noticeable importance is that this is primarily a visual representation of genetic evolution, and the fitness is observable through behaviour instead of numerical values making it more pleasing to the eye.

### 3 Experiments

#### 3.1 Experiment 1: “Evolving Shakespeare”

For our first experiment we used a genetic algorithm composed by Daniel Shiffman. The program uses DNA strings to compute the solution to a problem; the goal is to generate a specified sentence. In the original source code this was the string 'To be or not to be' but to speed up things we changed this into 'hello world'

A member of the population is a string of characters. Because we already know what our goal is, we can compare each member of the population with this eventual goal. For all members of the population we compare each position in its string with the corresponding position in the goal sentence. Finally we compute a percentage of 'matching' characters: this percentage will be our fitness.

Like most normal evolutionary algorithms, Shiffman's program has a strong preference for the fitter members of the population. These members are more likely to come up with the solution to the problem and are therefore allowed more often in our lottery of love, the mating pool.

Because they appear more often in the lottery, they are more likely to be selected as a dad or a mom for the next generation.

We had to find a way to somehow see the individual members live their insignificant little lives in service of the system. This is exactly what we implemented next.

The members of the population are represented by coloured circles. The colour of the circle represents the fitness of this members DNA: Black meaning a low fitness and white meaning a high fitness. When the program is run now the population as a whole quickly follows the path set out by its fittest members. When looking at the circles you see there's no members 'left behind' in darker colours. This is efficient but not very interesting from a serendipitous point of view.

One of the alternative approaches we wanted to try is crossbreeding the fittest member of the population with the least fit member. This would ensure a greater variety in the population (fig. 2) and thus make it easier for serendipity to occur. The algorithm does still select based upon the features of the members, it just allows a wider variety of flavours to co-exist. It is curious to notice the relative efficiency of this algorithm. The solution was found rather quickly considering there is just one couple of parents available. This is apparently caused by the greater range of interesting alternatives when selecting a lover from the lottery.

We realized the members of our population are not actually aware of the goal they are trying to achieve. By removing the advantage that cool individuals have in the lottery of love we tried to implement this in the algorithm, effectively removing “fitness” and not selection. We wondered if this would still lead to some evolution in a beforehand unforeseen way. There was barely any progress, the circles stayed mainly black with a few dark grey ‘waves’ once in a while.




Fig. 2 The result of our experiment: a more diverse population.

While conducting these experiments we took the opportunity to play some more with the settings of the algorithm to see how this would influence performance and behaviour. It's fun and fascinating to see the major influence of certain changes that seemed trivial when you made them. However: other changes didn't really yield the results we craved. Clearly, something was amiss here!

The problem of this approach is of course the limits of the algorithm. Because we only evaluate one specific goal we can't really tell whether we might have somehow achieved another, unknown goal at some point. The colours of our population only tell you how good they are at representing the solution to the specified goal. Since we still wanted to emulate evolution in a non-conventional way it was time to find more appropriate software.

### 3.2 Experiment 2: “Disney meets Darwin”

The second experiment was done using the “Disney meets Darwin” algorithm made by Jeffrey Ventrella. The goal was to find out interesting methods to evolve behaviour. Because the output is visual, it’s easy to detect subtleties.

The program allows customization of some features. One of them is the option to determine what behaviour should be rewarded, giving the virtual creature (fig. 3) a higher fitness function. Without any goal in life we expected the creatures to evolve some unpredictable behaviour. It would still be possible to reach specialized patterns, which are usually only generated through steering evolution that way, using the fitness function. Sadly though, using this technique leads to a predictable goal, while we are curious about patterns we did not expect, this is why we first generated some simulations where all goals were removed from the program. There were no rewards, better yet, no objectives.




Fig. 3 Some stick-like creatures from Ventrella’s simulations.

From this we found that earlier in evolution (in the first generations) the individuals exhibited a multitude of behaviours that were all different from one another. As evolution progressed, they tended to unify their behaviour, and latter generations produced the same behaviour over and over.

Another approach was to handpick individuals that were producing awkward behaviours (what the observer deemed interesting) and making this behaviour influence the next generation, persist in evolution. The behaviour preferred by the observer could not easily be put into terms of a fitness function. Still it had the same effect as using a fitness function, as future generations explored variations of this behaviour. The behaviour seemed more exploratory though, producing more deviations.

A follow-up experiment was conducted under the same circumstances, with handpicked individuals. This time though in each generation the behaviour chosen was the least close to the one selected on the previous generation. This attempt to ensure variation yielded interesting results, as now there were no fitness functions, but a multitude of behaviours that had nothing in common. This was the experiment that produced the biggest amount of variation.

By creating as much variation as possible, more and more individuals were discovered to act funny. Behaviour that would normally not occur, because evolution would destroy it by preventing these individuals to breed, was now flowering and presented a pleasing sight to behold.

## 4 Conclusion

The conclusion of this project is that though beauty might not lie in the eye of the beholder, this eye is certainly very important for observing something and deciding whether it is beautiful. Though we may have come up with some funny things, it is hard to decide whether something is ugly or beautiful, if there's no aesthetic aspect to it that can be recognized by us. The second experiment was therefore more successful than the first one, because at least the individuals showed observable behaviour.

Another conclusion is that the simulated evolution is not the best tool to study the phenomenon of evolution, because in nature there is no solution to a problem, there is just survival and extinction. A species might be very fit for its environment, but that doesn't mean it's the solution to a problem.

The panda is getting extinct, some people mourn this fact, but why should something that is evolutionary unsuccessful be preserved? This is because it has some aesthetic value to these panda-lovers. We hoped to find the virtual panda in our two genetic algorithms, but did not succeed. We couldn't prove its existence maybe, but it still might be out there!

## References

- [1] Minsky, Marvin. Negative Expertise. *International Journal of Expert Systems*. 1994, Vol. 7, No. 1, 13-19.
- [2] McLuhan, Marshal. *Medium is the Message*. Bantam book, 1967.
- [3] Stuart Russell en Peter Norvig (2003). *Artificial Intelligence: A Modern Approach*. Prentice Hall.
- [4] Daniel Shiffman's website: <http://shiffman.net/>
- [5] Jeffrey Ventrella, *Disney Meets Darwin; An Evolution-based Interface for Exploration and Design of Expressive Animated Behavior*, Medialab MIT, 1994.