Can the Red Queen Help Catch the Snark?
A Co-Evolutionary Waveform Transformation Approach

Marcelo Caetano[1], Jônatas Manzolli[2], Fernando Von Zuben[2]

[1]IRCAM – 1 place Igor Stravinsky – Paris, France 75004

[2]LBiC and NICS - University of Campinas - Brazil PO Box 6159

caetano@ircam.fr, vonzuben@dca.fee.unicamp.br, jonatas@nics.unicamp.br

Abstract. *Many works, ranging from biology to optimization techniques, claim to have encountered the Red Queen face to face**. *Nobody has ever been able to capture the Snark, though. In this paper, we describe a novel framework for waveform transformation entitled 'The Hunting of the Snark' based on the Red Queen Principle, which, in Lewis Carroll's words, states that "it takes all the running you can do, to keep in the same place". As an extension of our previous work that consisted of the application of evolutionary techniques to waveform synthesis, we generated waveform transformations that resulted from the exploration of the soundspace driven by a variation of the evolutionary 'arms race' paradigm applied to two populations of waveforms, denominated Predators and Prey. A mathematical distance between these populations gives fitness evaluation and the Predators chase the Prey, who try to escape. This spatial dynamics, which was dubbed 'Snark chase', gives rise to the parallel temporal evolution of two sets of waveforms that have the potential to be used in music composition, improvisation and possibly real-time performance. The genetic sound operators such as crossover, mutation and selection are explained in the light of the present implementation and preliminary results are shown.*

1. Introduction

Within the western musical framework, ever since the very dawn of electroacoustic music, the focus of a group of composers has been broadened from musical notes to musical sounds. Some compositional techniques sprouted from the manipulation of these sounds resulting, for example, in transformations of sonic features such as timbre or sound quality. Composers such as Trevor Wishart (2000) and Leigh Landy (1993) have produced works that can be said to gravitate around the utilization of sound transformations and its aesthetical implications. Early techniques included the use of analog equipment to store and manipulate the musical material (sounds themselves). An immediate consequence is that the results were restricted by the procedural limitations of the equipment itself. Enter the digital computer and its great musical potential and flexibility in sound manipulation. According to John Chowning (2000), digital waveforms can represent virtually any sound imaginable *given the correct sequence of numbers (digital sound samples)*. Max Mathews (1963) stated that "there are no theoretical limitations to the performance of the computer as a source of musical

* See, for example, Lythgoe 1998 and Paredis 1997.

sounds, in contrast to the performance of ordinary instruments”. Therefore, with an appropriate representation of all the possible sounds, denominated the soundspace, and a proper manipulation of certain properties of the representation and/or exploration of certain regions of the soundspace, the process of sound transformation corresponds to the generation and manipulation of sounds that evolve in time and therefore can be potentially applied to music composition, improvisation and real-time performance.

Sound transformations are usually associated with timbral metamorphosis [Landy 1993]. Since the pioneering work of Helmholtz (1885), timbre has been closely associated with spectral contents. Consequently, most sound transformation techniques described in the literature make use of manipulations in the spectral domain by means of a sound model that attempts to describe exhaustively the temporal evolution of the partials over the course of the sound. These initiatives rely on the analysis/synthesis paradigm, which requires the previous analysis of a sound to permit the synthesis step. Among the most famous such models are the Phase Vocoder [Flanagan 1966] and spectral modeling synthesis (SMS) [Serra 1998], [Smith 1997]. Moorer (1978) has extensively described how the phase vocoder can be used in computer music applications and Wishart (2000) and Landy (1993) have outlined applications of the Phase Vocoder in sound transformations in the perspective of the Composers’ Desktop Project (http://www.composersdesktop.com). There are approaches to both models attempting to calculate high-level descriptors [Serra 1998], [Amatriain 2002], [Tardieu 2004] that would allow us to manipulate salient timbral features of sounds independently, without affecting the other dimensions [Jensen 1999]. There is no consensus on what or how many these features are, however [Caetano 2005, 2007]. The approach most closely related to this work is that of Caetano (2007) based on the application of bio-inspired algorithms featuring characteristics of self-organization as tools to explore the soundspace of digital waveforms and generate sound transformations for music composition and improvisation. The method that will be presented in this paper is a variation of the evolutionary waveform synthesis (ESSynth) approach [Manzolli 2001a,b], which follows the principle that soundspace is virtually any digital waveform and that sounds can be generated by means that would promote the emergence of novel or even unexpected results. In other words, contrary to the analysis/synthesis approach, we propose the exploration of the soundspace of digital waveforms using evolutionary computation methods. Evolutionary computation is usually associated with creativity [Bentley 1999]. Dawkins (1986) stated that “as the search space gets larger, more and more sophisticated searching procedures become necessary. Effective searching procedures become, when the search space is *sufficiently* large, indistinguishable from true creativity”. There are computer models that illustrate well this argument and constitute an instructive bridge between human creative processes and the evolutionary creativity of natural selection [Latham 1992]. In this framework, sound transformations are evolving digital waveforms. Caetano *et al.* (2005) have proposed the evolution of a population (set) of waveforms towards a static target set. This process ends when (or if) the waveforms reach the target. The focus of the current investigation is to co-evolve both sets, named Predators and Prey, allowing the evolutionary process to run ideally indefinitely, which is also known in the literature as an arms race [Dawkins 1986]. The goal of the Predator set is to chase the Prey (minimize their distance to the Prey set), while the Prey must escape the Predators (maximize their distance to the Predator set).

The next section briefly overviews the red queen principle (RQP) and its connection with co-evolution, attempting to present the motivation of the use of co-evolution as a paradigm for sound transformation. We then proceed to thoroughly explain 'The Hunting of the Snark'. The algorithm, the genetic parameters and operators are properly sketched. Next, we focus on the results, comprising the experiment and the discussion of its outcome. Finally, the conclusions and future work are presented.

2. The Red Queen Principle in Co-Evolution

Evolution can be defined as the accumulation of change in the individuals of a population over the generations due to natural selection. Holland (1975) devoted himself to the study of adaptive natural systems and, inspired by biological evolution, he proposed Genetic Algorithms (GAs) to indicate that adaptation mechanisms can be properly implemented in computers. GAs mimic nature in accordance with Darwin's survival of the fittest principle, exchanging information in a structured yet random way. GAs codify attributes that fully characterize the elements of a search space, using the language of computers. The resulting search space contains the candidate solutions, and the evolutionary operators will implement exploration and exploitation of the search space aiming at finding global optima. The GA iteratively manipulates populations of individuals at a given generation by means of simple genetic operations of selection, crossover and mutation. GAs tend to lose variability over the generations [Davis 1991]. Co-evolution helps preserve and even generate variability [Palazios 2004]. This can be understood as different strategies for catching the prey (or, alternatively, for escaping) being developed by individuals over the generations. Co-evolution refers to the simultaneous evolution of two or more interacting species. This framework for evolution may be conceived so that the species cooperate or compete. The application of competitive co-evolution to problem solving has been of interest in the GA community because *competition*, in its most general sense, encourages the generation of better *competitors* [Paredis 1995]. To the best of our knowledge, there have been no proposals for the use of co-evolution in waveform transformation.

2.1. Evolutionary Arms Races and the Red Queen Principle

There are ways in which mutation and natural selection together can lead, given enough time, to a building up of complexity, beauty and efficiency of design. Evolutionary 'arms races' is one of them, under the concept of co-evolution. There are arms races between predator and prey, parasites and hosts, even between males and females within one species [Dawkins 1986]. Arms races consist of the improvement in one lineage's (say prey animal's) equipment to *survive*, as a direct consequence of improvement in another (say predator's) lineage's evolving equipment [Dawkins 1986]. The principle of zero change in success *rate*, no matter how great the evolutionary progress in *equipment*, has been given the memorable name of the Red Queen Principle (RQP) by the biologist Leigh van Valen (1973). In *Through the Looking Glass* [Carroll 1872] the Red Queen seized Alice by the hand and dragged her, faster and faster, on a frenzied run through the countryside, but no matter how fast they ran they always stayed in the same place. Alice was understandably puzzled, saying, 'Well in our country you'd generally get to somewhere else-if you ran very fast for a long time as we've been doing.' 'A slow sort of country!' said the Queen. 'Now, here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!' The RQP is directly related to positive feedback through symbiotic relations between individuals of the different populations competing, it raises the question "Can a natural evolutionary system support continual change?"

3. The Hunting of the Snark

The 'Hunting of the Snark' method was named after Lewis Carroll's homonymous nonsense poem that describes the voyage of a crew after an imaginary creature that cannot be caught. In this paper, 'The Hunting of the Snark' has two populations (sets) of sounds, Predators and Prey. Each population contains a number of individuals that are sounds represented by N samples of a digital waveform at a sampling frequency of FS samples per second. The waveform is the genetic code (genotype) that carries all the information regarding the sound and can be manipulated. The resultant sound (phenotype) is the characteristic that can be perceived. The 'Hunting of the Snark' algorithm is shown in Figure 1 and the input parameters are summarized in Table 1 and were adapted from Palacios-Durazo and Valenzuela-Rendón (2004). Fitness evaluation is given by a distance measure (Hausdorff metric) between Predators and Prey as follows: we calculate the Euclidean distance between two waveforms regarding them as vectors with N dimensions. That is, each sound sample is considered as a vector component. As shown in the algorithm in Figure 1, the distance from each of the K Predators to all the M Prey is measured, resulting in a K by M matrix of distances. For each Predator, fitness is the minimum distance found from all the M Prey. The same is done for each Prey. Now that each individual has a distance value attributed, the best Predator is the closest to the Prey, and the best Prey is the furthest away. Every individual "sees" all other individuals and diversity is not necessarily preserved in terms of speciation. This particular spatial dynamics described in our work was dubbed 'Snark chase'. The fitness measure produces a non-linear mapping of the genotype into the phenotype and the 'Snark chase' dynamics conducts the exploration of the soundspace. Thus, we defined the 'Snark chase' as "the improvement in the Prey's evolving equipment to *escape*, as a direct consequence of improvement in the Predator's evolving equipment to *chase* the Prey."

```
(*Initialize Populations*)
        Load predefined Predator population
        Load predefined Prey population
(*Main Cycle*) generations
    repeat
(*Competition Cycle*) Snark chase
    for each p1 ∈ Predators
    for each p2 ∈ Prey
        measure distance between p1 and p2
    end for
    end for
        Fitness of Predators and Prey calculated based on the minimum distance of each individual to all individuals
        in the competing population
            One generation of a GA is applied to Predators
                Crossover, mutation, selection
            One generation of a GA is applied to Prey
                Crossover, mutation, selection
    until termination criteria met
```

Figure 1. 'The Hunting of the Snark' algorithm.

3.1. Genetic Parameters

The user will adjust the search according to predefined requirements achieved by the manipulation of the parameters of the genetic algorithm. It follows closely the original ESSynth method [Manzolli 2001a,b], [Caetano 2005].

3.1.1. Size of the Population

The size of the population directly affects the efficiency of the GA [Davis 1991]. A small population supplies a small covering of the search space of the problem. A vast population generally prevents premature convergences to local solutions. However, greater computational resources are necessary [Davis 1991]. In 'The Hunting of the Snark', the size of the population does not change along the generations.

3.1.2. Coefficient of Mutation

It determines the probability of mutation [Holland 1975]. A properly defined coefficient of mutation prevents a given individual of the population from stagnating in a particular position and also promotes exploration and exploitation of the search space. A very high coefficient of mutation causes the search to become essentially random and increases the possibility of destroying a good solution [Davis 1991]. In 'The Hunting of the Snark', the coefficient of mutation ranges from 0 to 1. It is user defined (input argument to the algorithm) and does not change along the generations, being responsible for limiting the noisy distortion caused by the mutation operator, described in section 4.2.2.

3.2. Genetic Operators

It is the genetic operators that transform the population along successive generations, being responsible for the emergence of evolution in computers. A standard genetic algorithm evolves, in its successive generations, by means of three basic operators, crossover, mutation and selection, described as follows.

3.2.1. Crossover

It represents the mating between individuals [Holland 1975]. The central idea of crossover is the propagation of the characteristics of the individuals in the population by means of the exchange of information segments between them, which will give rise to new individuals. In 'The Hunting of the Snark', crossover operation exchanges chromossome segments, i.e. a certain number of samples, between individuals sharing the same 'gene pool' through a smooth transition. There is no crossbreeding between Predators and Prey. There are many different possibilities concerning the selection of which individuals will be crossed-over with each other. One must be aware, though, that

Table 1. Input parameters of 'The Hunting of the Snark'

N	Number of samples per individual
FS	Sampling rate
K	Number of Predators
M	Number of Prey
NumInt	Number of generations
coefMut	Coefficient of mutation

the selection operator for crossover has a great influence in the diversity of the genetic pool. We wish to preserve diversity to its maximum to guarantee a good exploration of the search space along the generations. However, diversity maintenance must not slow down convergence. Many strategies were tried beforehand: randomly selecting both individuals, crossing each individual with a randomly chosen individual, crossing the best individual with a randomly chosen individual. A suitable strategy was found to be crossing each individual with the best individual in each generation to pass along the best individual's gene pool while maintaining diversity.

The crossover operation is depicted in Figure 2 and consists in adding a sound segment from the best individual to each individual of a given population. The operation is somehow similar to granular synthesis and can be interpreted as adding a single sonic quantum from the best individual. The segment is obtained as follows: we randomly generate the width of the segment, called *slice*, between *N/10* and *N/5*. Next we generate the center of the segment, which must be a value between *1+slice/2* and *N-slice/2* to guarantee that the entire segment always falls between *1* and *N*. We then proceed to create the crossover functions, shown in Figure 3 parts a) and c). The crossover function in part c) of Figure 3 is a Hanning window centered at *center* with width equal to *slice* and zero-padded to size *N*. The crossover function in part a) of Figure 3 is one minus the previous function. Then the function in part a) of Figure 3 is applied to one individual of a given population, resulting in the segment seen in Figure 3 part b), and the crossover function shown in Figure 3 part c) is applied to the best individual of the same population, resulting in the segment seen in Figure 3 part d). The resultant segments are added producing another individual, denominated Breed, shown in part c) of Figure 2. It should be noted that the individual being crossed-over with the best individual in that generation is preserved almost entirely, receiving only a narrow segment from the best. The crossover functions are generated for each individual of both populations.

3.2.2. Mutation

It produces random modifications and is responsible for the introduction and maintenance of genetic diversity in the population [Holland 1975]. In 'The Hunting of the Snark', the mutation operator is as follows: *N* values randomly generated between 1−*coefMut* and 1 are used to multiply each of the corresponding elements of a given individual (waveform) of the population. This operator introduces a certain noisy distortion to the original waveform that is equivalent to a non-linear perturbation in the genotype that reflects in phenotype. Notice that the higher the value of *coefMut*, the stronger the perturbation.

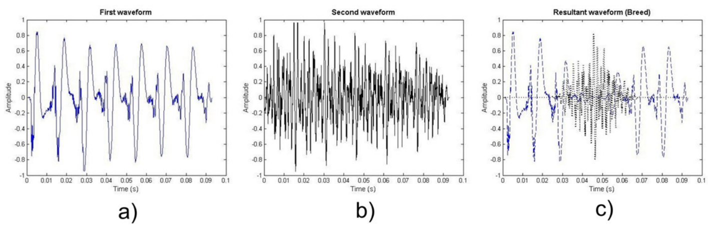

Figure 2. Depiction of the crossover operation in time domain. Part a) and b) show the waveforms to be crossed-over and part c) shows the result of the operation.

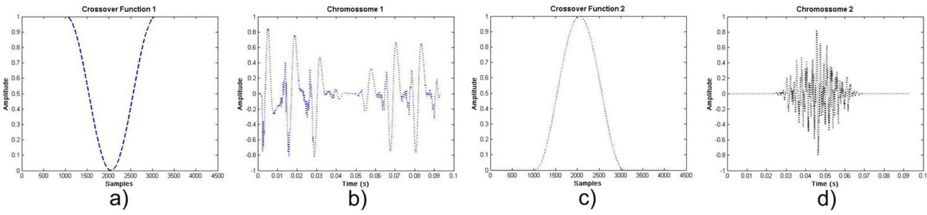

Figure 3. Crossover functions and the sound segments that result after their application.

3.2.3. Selection

After mutation all individuals are passed to the next generation. The absence of selection for the next generation is a consequence of practical experience. We found that all the selection strategies we tried (roulette wheel, random selection and rank) caused too much selective pressure and consequently the diversity dropped to zero very soon.

4. Results

The purpose of this section is not only to present some outcomes of co-evolution in terms of sonic results, but also to allow the reader to have a better understanding of how it works. For such, we need to highlight that the evolution of the waveforms corresponds to a sound transformation procedure, that the Snark chase dynamics is responsible for how this transformation is done and that co-evolving the sounds has the desired side effect of helping the preservation of diversity, apart from the obvious fact that co-evolution in this case means that both populations of waveforms are being transformed at the same time. The result of this experiment can be heard at `ftp://ftp.dca.fee.unicamp.br/pub/docs/vonzuben/SBCM2007/sound_samples.zip.`

The overall result of the application of the method can be easily inferred as a sound transformation procedure that adds partials (spectral complexity) to the sounds over the generations mainly because of the crossover operator. In order to try to "show" the transformation procedure, the output sound set resulting from a run of the program will be shown and discussed. The Predator waveforms were taken from a recording of the electric bass and the Prey from piano chord sounds. The resultant waveforms of one individual from each population will be presented in generations 1, 2, 5, 10, 50 and 100 to illustrate the transformations of the waveform along the generations. The generational waveform display attempts to draw attention to the variability each individual presents over the generations, despite the fact that it preserves some of its original information contents. A 3D plot of the short-time Fourier transform (STFT) – thereon referred to as dynamic spectrum - of one individual from each of the populations in the first and last generations will be compared to exemplify the spectral transformations induced by the method. The dynamic spectrum of the original and resultant individual highlights the changes to the temporal envelope evolution of the partials induced by the method. Finally, we show plots of the distances of the individuals along the generations. It might be surprising at first to see the distance value being used as a measure of diversity since different individuals can have the same distance. On the other hand, one individual can only have a single distance value. Therefore, different distance values necessarily mean different individuals. We can even go one step further and suggest that small distances imply small differences (and viceversa) because waveforms (individuals) that are only slightly different present a small distance (they are correlated vectors). But we must be careful because, although it is correct for the genotype (waveforms), we know it is far from necessarily true for the phenotypic differences (how they sound). All we need to remember is that differences in phase do not affect greatly the sound but alter considerably the shape of the waveform.

4.1. Experiment

Since the experiment basically consists of a run of the program, we are committed to show that the result is a different sound transformation procedure for each population. Due to the loss of variability over the generations, the individuals of the same population do tend to sound alike after one or two hundred generations. However, we know that the initial individuals were different, so we will always have variations of the transformation. Therefore, we consider the results satisfactory if we verify that one individual from each population actually corresponds to a sound transformation and that they are different. The genetic parameters adopted in the experiment were 10 individuals in both the Predator and Prey populations to supply a satisfactory coverage of the search space. The program was run for 200 generations. *FS* is 44,100 so the highest representable frequency is 22,050 Hz. *N* was set 4096 so each chromosome is a wave-format sound segment of approximately 0.0929s. The coefficient of mutation was set 0.05 and was obtained empirically because higher values have shown to distort so much the waveforms that the results were almost too noisy and masked the transformation. Next, we present the waveforms of both Predator 1 in Figure 4 and Prey 1 in Figure 5 in generations 1, 2, 5, 10, 50 and 100. The top row from left to right shows generations 1, 2 and 5 and the bottom row, also from left to right, show generations 10, 50 and 100. Figure 6 shows 3D plots of the STFT of the original Predator 1 and Prey 1 and their transformed versions after 200 generations. The top row shows Predator 1, from left to right in the first and in the last generation. The bottom row shows Prey 1, from left to right in the first and in the last generation. These figures are similar to spectrograms in that they show the temporal evolution of the envelope of each partial frequency along the course of the sound. Finally, Figure 7 shows plots of the fitness values of both populations in all 200 generations. Part a) shows fitness value of all 10 Predators in all 200 generations, part b) shows fitness value of all 10 Prey in all 200 generations and part c)

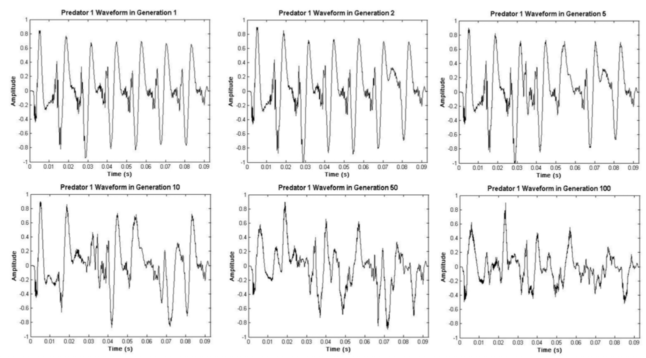

Figure 4. Waveform evolution for Predator 1. From left to right, the top row shows the waveform Predator 1 in generations 1, 2 and 5 and the bottom row in generations 10, 50 and 100.

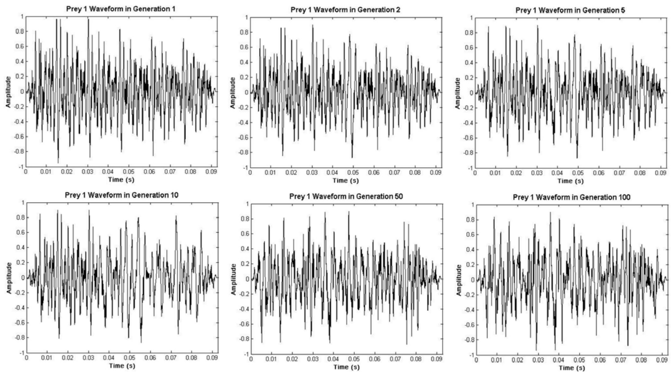

Figure 5. Waveform evolution for Prey 1. From left to right, the top row shows the waveform Prey 1 in generations 1, 2 and 5 and the bottom row in generations 10, 50 and 100.

on top fitness value of the best Predator in each generation and at the bottom fitness value of the best Prey for all 200 generations.

4.2. Discussion

Figures 4 and 5 show clearly the waveform transformation in the course of the generations. For example, from a close examination of Figure 4, it becomes clear how in the second generation Predator 1 already possesses one cycle that is different from the others. It comes from the best individual in that generation via the crossover operator. The Prey 1 waveform shown in Figure 5 is somewhat noisier than that of Predator 1, making it more difficult to visually identify the changes it suffers over the generations. However, it becomes more evident if we examine the dynamic spectrum from Prey 1 in the first and last generations in Figure 6, which illustrates the adding up of spectral complexity in the individuals along generations. The temporal envelopes of the partials have changed a great deal, which means that not only the waveform itself has changed, but also the spectral information. We need to bear in mind that while changes to the

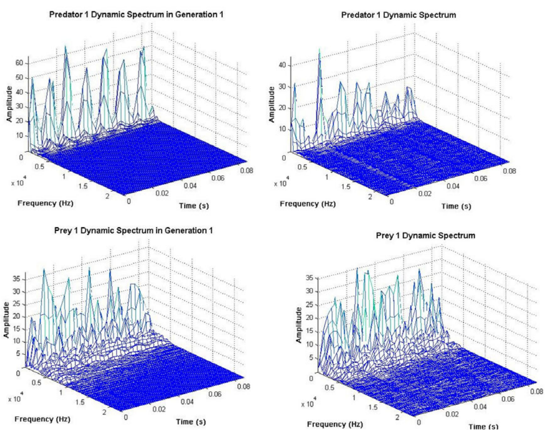

Figure 6. Spectro-temporal transformation to Predators and Prey. The top row shows the dynamic spectrum for Predator 1 and the bottom row for Prey 1. The left column shows both individuals in the first generation and the right in the last generation.

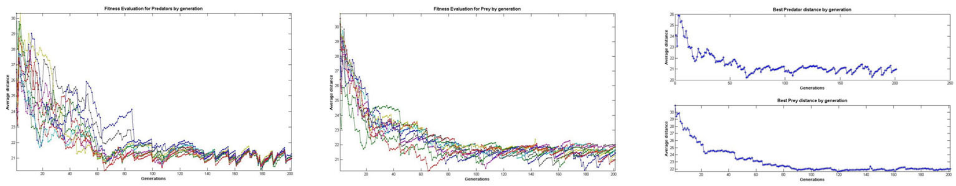

a) b) c)

Figure 7. Distance evolution along generations. Part a) shows the distance value for all the Predators, part b) for all the Prey and part c) shows the distance value of the best Predator on top and best Prey at the bottom.

phases of the partials modify the waveform but are not perceived in general, changes to the amplitude of the partials also affect the shape of the waveform and are certainly perceived as a different sound. Obviously, the same goes for Predator 1 also shown in Figure 6, which permits us as well to "see" the result of the distortion caused by the mutation operator as the noisy contents in high frequencies that the resultant waveforms present and that were not present in the dynamic spectra of the original waveforms.

Finally, to analyze Figure 7, it is important to remember that the waveforms are kept with the same label along all the generations. This means that the waveform that was loaded as Predator 1, for instance, is never relabeled during the process; it only suffers the changes made by the crossover and mutation operators. Furthermore, we expected the overall distance of Predators to decrease, indicating that they do chase the Prey. This can be confirmed by visual inspection of Figure 7 part a), which shows a decrease of distance values over the generations that means that the Predators do get closer to the Prey by means of the genetic operators alone. The evolution of the distance value for the best individual in each generation confirms this general tendency. Figure 7 also hints at the maintenance of diversity because of the different fitness values that are represented by the curves farther apart. For the Predators, from the beginning of the generations until at least halfway through to the end there is considerable variability, but then around the 150th generation the Predators seem to converge to a single oscillating genotype (waveform) that is supposedly the most adapted to chase those specific Prey. This oscillating behavior seems to confirm the co-evolutionary predator-prey dynamics. The Prey do not group together like this over the generations, keeping roughly the same diversity. This might reflect a strategy developed by the Prey to keep away. Maybe keeping a more diverse gene pool helps the Prey avoid be reached by the Predators. Although it is too soon to jump to conclusions like this, we would like to be able to find out whether a specific genotype for the Prey induces the Predators to converge to a certain specific region of the soundspace similarly to what the static target sounds in ESSynth cause (Caetano et al. 2005). This would mean we could "direct" the transformations of the Predators towards some desired general result with the right population of Prey.

5. Conclusions

We have described 'The Hunting of the Snark', a waveform transformation method that has co-evolution as paradigm. It can be seen as a variation of the evolutionary sound synthesis method (ESSynth), which applies a GA to a population of waveforms that are driven towards another static population as target. Here, the target waveforms move away motivated by a variation of a model of co-evolution. We denominated the two populations of waveforms Predators and Prey and the sound transformation method consists of the Predators chasing the Prey and the Prey trying to run away from the Predators. Both populations are being transformed in parallel in what we have dubbed Snark chase. Ideally, the Snark chase dynamics means that, as long as we have variability, the process goes on. Put in a different way, the Predators will never reach the Prey if they keep successfully running away, but they will keep chasing them anyway. This supposedly means that the sound transformation process is virtually endless because the target is moving away. However, due to the nature of the crossover operator, its exhaustive application over hundreds of generations forces homogeneity in the genotype (waveforms), which necessarily means that there will be a loss of diversity along the way. This loss of diversity corresponds to the convergence of the algorithm. We expect, though, the Predators to try and catch the Prey and the Prey to run away as fast as they can. This fact should be reflected as a decrease in distance values for Predators. We have indeed found a decrease in distance for the Predators over the generations that can be interpreted as the Predator population approaching the Prey. Also, upon convergence, there is indication of oscillatory behavior, characteristic of co-evolution. The Prey are being transformed just like the Predators. The very same crossover, mutation and selection operators are applied to both populations and the waveforms of one individual from each population in some predefined generations we have shown confirm that both Predators and Prey are being transformed, endorsing the use of co-evolution as paradigm for evolutionary sound transformation.

Future work might include testing different mutation operators; comparing the results of evolutionary and co-evolutionary waveform transformation, verifying if a certain group of Prey drives the Predators towards a specific region of the soundspace and even experimenting with multi objectives, so that the Predators would not simply chase the Prey, corresponding to different kinds of competition.

Acknowledgements

This work was fully developed under the supervision of Profs. Fernando Von Zuben and Jônatas Manzolli both at the University of Campinas and supported by FAPESP (process 03/11122-8) and CNPq. The main author is presently supported by CAPES (process 4082-05-2) and advised by Prof. Xavier Rodet at IRCAM.

References

Amatriain, X., Bonada, J., Loscos, A., Serra, X. ``Spectral processing," in DAFX - Digital Audio Effects: 373-438, John Wiley and Sons, 2002.

Bentley, P. (1999). Evolutionary Design by Computers. S. F., Morgan Kaufmann.

Caetano, M., Manzolli, J. and von Zuben, F. Self-Organizing Bio-inspired Sound Transformation. EVOMusart workshop. Proc. of the EVOStar, 2007.

Caetano, M., Manzolli, J., von Zuben, F. Interactive Control of Evolution Applied to Sound Synthesis. Proc. of FLAIRS, EUA, 2005.

Carroll, L. Through the looking glass and what Alice found there. Macmillan, 1872.

Chowning, J. M. Digital Sound Synthesis, Acoustics and Perception: A Rich Intersection. Proc. of the DAFX-00, 2000.

Davis, L. “Hanbook of Genetic Algorithms”. New York:Van Nostraud Reinhold,1991.

Dawkins, R., “The Blind Watchmaker”, Penguin Books, 1986.

Flanagan, J., Golden, R. Phase Vocoder, Bell Syst. Tech. J. 45:1493-1509, Nov. 1966.

Helmholtz, H. von. On the Sensations of Tone. London, Longman, 1885.

Holland, J., Adaptation in Natural and Artificial Systems. Univ. of Michigan Press, 1975.

Jensen, K. Timbre models of musical sounds. Ph.D. Dissertation, Dep. Computer Science, Univ. Copenhagen, 1999.

Landy, L. (1993) Sound Transformations in Electroacoustic Music. http://www.composersdesktop.com/landyeam.htm. Access in 06/2007.

Latham, W., Todd, S., “Evolutionary Art and Computers”, Academic Press, 1992.

Lythgoe, K. A., Read, A. F. Catching the Red Queen? The advice of the rose. Trends Ecol. Evol. 13: 473-474, 1998.

Manzolli, J, A. Maia Jr., Fornari J.E. and Damiani,F. 2001a. The Evolutionary Sound Synthesis Method. ACM Multimedia, Ottawa, Ontario, Canada.

Manzolli, J, A. Maia Jr., Fornari J.E. and Damiani,F. 2001b. Waveform Synthesis using Evolutionary Computation. Proc. of the VII SBCM.

Mathews, M. V. The Digital Computer as a Musical Instrument. Science, 142 (3592), pp. 553-557, 1963.

Moorer J. The use of the phase vocoder in computer music applications, Journ. of the Audio Eng. Soc., 26 (1/2), 1978.

Palacios-Durazo, R., Valenzuela-Rendón, M. Similarities between Co-Evolution and Learning Classifier Systems and their Applications, GECCO 1: 561-572, 2004.

Paredis, J. Coevolutionary computation. Art. Life. 2 (4): 355-375, 1995.

Paredis, J. Coevolving cellular automata: be aware of the red queen! Proc. of ICGA: 393-400, 1997.

Serra, X. J. Bonada. "Sound Transformations Based on the SMS High Level Attributes", Proc. of DAFX Workshop, 1998.

Smith, J. O. (2007) Spectral Audio Signal Processing. CCRMA, Stanford University. http://ccrma.stanford.edu/~jos/sm/Spectral_Modeling_Synthesis.html. Access in 06/2007.

Tardieu, D. Synthèse et transformation sonore par descripteurs de haut-niveau. Master thesis, Université Aix Marseille, 2004.

Van Valen, L. A new evolutionary law. Evolutionary Theory, 1: 1-30, 1973.

Wishart, T. (2000) Computer Sound Transformation: A Personal Perspective from the UK. http://www.trevorwishart.co.uk/transformation.html. Access in 06/2007.