iSoundScape: Adaptive Walk on a Fitness Soundscape

Reiji Suzuki[1,2], Souichiro Yamaguchi[1], Martin L. Cody[2], Charles E. Taylor[2], and Takaya Arita[1]

[1] Graduate School of Information Science / SIS, Nagoya University
Furo-cho, Chikusa-ku, Nagoya 464-8601, Japan
[2] Department of Ecology and Evolutionary Biology
University of California
Los Angeles, CA, 90095, USA

Abstract. Adaptive walk on a fitness soundscape [7] is a new kind of interactive evolutionary computation for musical works. This system provides a virtual two-dimensional grid called a "soundscape" in which each point corresponds to a genotype that generates a sound environment. By using the human abilities of localization and selective listening, the user can "walk" toward genotypes that generate more favorable sounds. This corresponds to a hill-climbing process on the "fitness soundscape." This environment can be realized by multiple speakers or a headphone creating "surround sound." In this work we describe two new applications of adaptive walk. The first is developed for creating spatially grounded musical pieces as an interactive art based on fitness soundscapes. The second provides a new way to explore the ecology and evolution of bird songs, from scientific and educational viewpoints, by exploring the ecological space of "nature's music", produced by populations of virtual songbirds.

Keywords: interactive evolutionary computation, musical composition, fitness landscape, surround sound, birdsongs, artificial life.

1 Introduction

Interactive evolutionary computation (IEC) has been used for optimizing a variety of artifacts which cannot be evaluated mechanically or automatically [8]. Based on subjective evaluations by a human, one's favorite artifacts in the population are selected as "parents" for new artifacts in the next generation. By iterating this process, one can obtain better artifacts without constructing them directly.

IEC has found use in a variety of artistic fields, including visual displays [6], musical compositions [9] and sound design [3]. While promising, IEC for musical composition has shortcomings because it is difficult to evaluate candidate sounds when a large number of them are played at once. Consequently, the users typically had to listen to each piece separately and evaluated them, one by one. This sequential evaluation of individuals has two costs: an increase in the total

C. Di Chio et al. (Eds.): EvoApplications 2011, Part II, LNCS 6625, pp. 404–413, 2011.

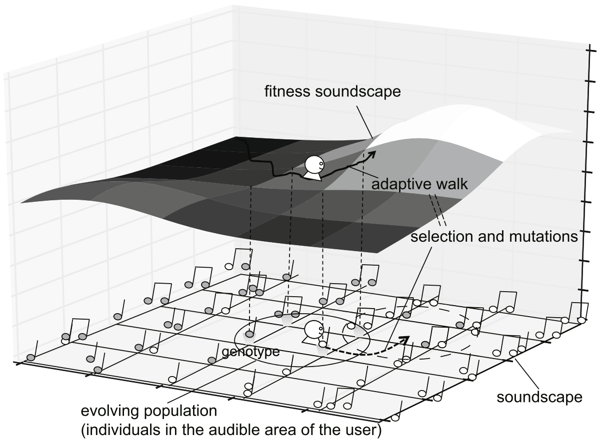

Fig. 1. A basic concept of adaptive walk on a fitness soundscape and its relationship with evolutionary computation

evaluation time (temporal cost), and the need for remembering the features of individuals in order to compare them (cognitive cost). Thus, it is necessary to limit the population of candidate pieces in each generation to a small size in order to decrease these costs, leading to inefficient selection and introducing substantial random genetic drift [1].

Recently, several studies proposed and developed surround-sound-based browsing or exploration of musical collections or sounds [5,4]. Also, to increase efficiency of IEC for sounds, Suzuki and Arita [7] proposed a new kind of IEC for musical works which was inspired by a biological metaphor, adaptive walk on fitness landscapes. A fitness landscape is used to visualize and intuitively understand the evolutionary dynamics of a population [10]. Recognizing that similar sounds could be placed nearby one another in a virtual landscape, they constructed a system that combined human abilities for localization and selective listening of sounds with a hill-climbing process on fitness soundscapes. This system enables a user to explore her favorite musical works by moving through a virtual landscape of sounds. Fig. 1 shows a conceptual image of this system. It may be summarized as follows:

– We assume a set of genotypes that can describe all possible musical pieces to be explored.
– We also assume a two-dimensional grid, and map every genotype to a unique grid point so that there is a correlation between the distance and similarity among genotypes on the grid. That is to say, similar genotypes are located nearby, while less similar ones are more distant in this grid.
– Each genotype can play its own musical pieces at its corresponding location in the grid. The resulting two-dimensional acoustic space is called “soundscape”, as shown in Fig. 1 (bottom).

- A user of this system, whom we will call the listener, has a location in the soundscape, and can hear the sounds of neighboring genotypes at the same time, if desired. The sounds come from different directions that correspond to their locations on the soundscape. This virtual environment can be realized by a multiple speaker system creating “surround sound.”
- Humans have a sophisticated ability to localize the direction of sounds, and can focus their attention in one direction or another. This is sometimes called the *cocktail party effect* [2]. By using their ability to localize, the listener can evaluate neighboring sounds at the same time. Their evaluation of goodness of the sounds gives each direction an analog to a gradient of fitness, thereby defining a third dimension – which may be thought of as a fitness surface on the soundscape – a “fitness soundscape.” The actual shape of the fitness soundscape will be different among listeners depending on their subjective impression, and can also change dynamically even for a single user.
- The listener is able to change her location – “walk” – along her own fitness soundscape by repeatedly moving toward the direction with an increasing fitness on the surface – i.e. in the direction from which more favored sounds are coming. This corresponds to the evolutionary process of the population in standard IECs. In other words, adaptive evolution of the population can be represented by a walk along the soundscape, determined by a hill-climbing process on the fitness soundscape.

This system can be regarded as a kind of evolutionary computation in the following sense: An evolving population in evolutionary computation corresponds to the set of genotypes on the soundscape whose sounds can reach the user. The movement of the user toward her favorable sounds corresponds to selection and mutation operations because less favorable genotypes disappear and new, slightly, different genotypes appear in the new population due to the shift of the audible area of the user. Although a crossover or recombination operation is not incorporated into the conceptual model, changing the scale and shape of the soundscape implemented in the prototype can contribute to maintaining the genetic diversity of the population (explained later).

Suzuki and Arita [7] constructed a prototype of the system using a personal computer to control a multi-channel home theater system with 7 speakers. They confirmed that listeners were able to search for their subjectively more favorable pieces by using this system. Their experience led to a proposal for improving evolutionary search in that system. Finally, they noted that a searching process on the soundscape was itself a new experience for the listener, suggesting that the system might be implemented as an art installation.

In this paper, we report on two variant applications of this concept. The first is developed for creating more spatially grounded musical pieces as a kind of interactive art, exploring better use of the surround sound environments with an iPhone or iPad. The second variant is to provide a new way to experiment with the ecology and evolution of birdsongs from both scientific and educational viewpoints — in a sense to “nature’s music” by songbirds.

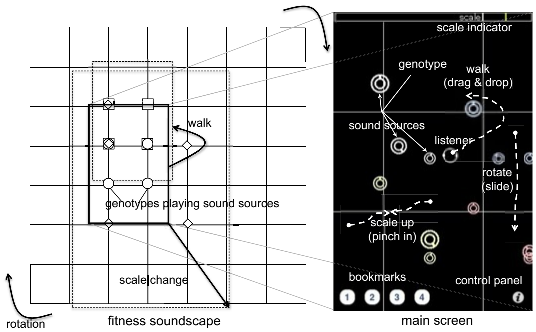

Fig. 2. Basic operation of the iSoundScape App (left) and a snapshot (right) of the main screen, with labels described in the text

iPhones and iPads provide a variety of sensors and interfaces, and have a rich software development library for sound works, such as Core Audio and OpenAL. The devices are now becoming standard platforms for applications of interactive arts based on sounds. For example, Brian Eno, a pioneer of ambient music, recently released a kind of interactive music box called "Bloom", in which a user can create drones playing periodically by touching a screen. Bloom and several variants are available at http://www.generativemusic.com/. In the applications described here, we use only a stereo (two-speaker) audio environment, rather than the richer environment provided by the 7-speakers in earlier studies. Something is lost, certainly, but we are finding that the intuitive operation and quick responses of a multi-touch interface may at the same time bring new benefits for exploring fitness soundscapes. The availability of world-wide channels to distribute applications from AppStore and to obtain feed-back from users are also benefits from this platform.

2 iSoundScape

We constructed an application for iPhone termed iSoundScape based on an adaptive walk on the fitness soundscape illustrated in Fig. 2. It was developed using Objective-C++ with XCode, and enlists the OpenAL API for creating surrounded sounds. It is currently available from the Apple App store at no cost.

Fig. 2 (left) shows a global image of a soundscape represented as a two-dimensional and toroidal grid space in this system. Fig. 2 (right) is a snapshot of a main screen of the implementation of an iPhone. It shows a part of the whole soundscape from a top view. An open circle in the center corresponds to the listener and generates her sound environment there. There are also several double circles that represent the sound sources mapped from genotypes on nearby grid points. The listener is assumed to face upward, and can hear the sounds emanating from the corresponding sources, in a manner described below.

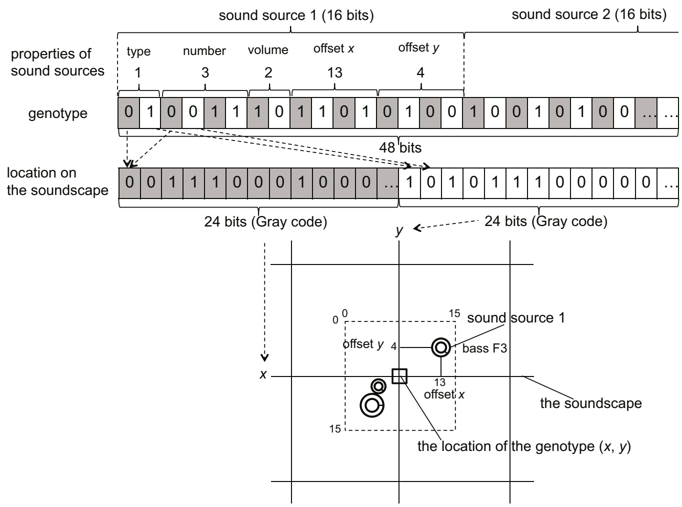

Fig. 3. The relationship between a genotype and its location on the soundscape, and one mapping from a genotype to a musical piece

2.1 Genotypes for Creating a Soundscape

Spatially grounded representations of musical pieces are utilized in this work, improving the use of the acoustic space of the soundscape compared to the simple scores often used in previous research [9,7].

Fig. 3 shows the relationship between a genotype and its location on the soundscape, as well as one mapping from a genotype to its musical piece. Each genotype is composed of a bit string with a length of 48, which contains enough information to code for more than 67 billion possible genotypes. We can imagine each 48-bit string as a point in a $2^{24} \times 2^{24}$ soundscape grid. The information of a genotype is used to determine its location on the soundscape, and also used to determine the properties of its corresponding musical piece as described below.

In order to determine the location on the soundscape on the $2^{24} \times 2^{24}$ grid, the bit string of a genotype is split into the two substrings composed of odd or even bits in the original genotype respectively, as shown in Fig. 3. Each substring is mapped to an integer value by using a Gray code (Reflected Binary Code), and the set of integer values is defined as the x and y location of the genotype on the soundscape. A Gray code maps each number in the sequence of integers $\{0, 1, \cdots, 2^{24}-1\}$ as a bit string with the length 24 so that consecutive numbers have binary representations that each differ by one bit. In this way, genotypes are closely correlated to their locations on the soundscape.

Each genotype generates and plays a musical piece distinctive to its location on the soundscape, composed of sounds from the three closest sound sources played on an infinite loop. In this prototype, the four nearest genotypes displayed on the screen generate the sound sources. In order to determine the property of these sound sources, the bit string of the genotype is also split into the three

Table 1. The sound clips

ID	type	unique number
0–15	piano (0)	C3–C5 (0–14) and no sound (15)
16–31	bass (1)	C3–C5 (16–30) and no sound (31)
32–47	synthesizer (2)	C3–C5 (32–46) and no sound (47)
48–63	drums / sound effects (3)	drums (48–55), singing birds (56, 57), sound of water (58, 62), call of goat (59), laughter (60), clock (61) and ambient sound at lobby (63)

substrings with the length 16, as shown in Fig. 3. Each substring determines a type of sound clip to play (4 types), a unique number in the type (16 different clips), a volume (4 levels), and a relative position from the grid point of the genotype (16×16 locations) on the soundscape, as shown in Fig. 3. The type of sound and the unique number in the type specifies one of 64 different sound clips (in wave format) as listed in Table 1. Piano, bass and synthesizer are each characterized as a short drone with a unique pitch. We also prepared some beats on drums and sound effects from various natural or social environments. Each sound source plays its sound clip with the specified volume repeatedly. Either even or odd bits are used from a genotype to determine the appropriate x or y location by relating each x and y value to the overall properties of the sound sources. Each sound source is represented as a double circle with a tab that indicates the ID of the sound clip. The size of the circle represents the volume of the sound source.

2.2 Operations for Adaptive Walk on a Fitness Soundscape

There are several basic operations to walk on the soundscape as follows:

Adaptive walk. After evaluating the musical pieces generated by the nearest four genotypes, a listener can change its location on the soundscape by dragging its icon. The directions of sounds coming from sound sources change dynamically according to the relative position of the listener during the movement. If the icon is moved outside of the central area of the grid, the soundscape scrolls by one unit of the grid as shown in Fig. 2 (left). The sound sources of the next four neighboring genotypes are then displayed and begin to play.

Rotation. By swiping a finger outside of the center area in a clockwise or counterclockwise direction, a listener can rotate the whole soundscape by 90 degrees (Fig. 2 left). By changing the orientation of the surrounding sound sources in this way, a listener can utilize the right / left stereo environment of the iPhone for localizing sounds that were previously in front of / behind the listener.

Scale change of the soundscape. By pinching in or out any place on the screen, a listener can decrease or increase the scale of the soundscape. It changes the Hamming distance between the nearest neighboring genotypes as shown in Fig. 2 (left) by skipping the closest genotypes and selecting distant genotypes. The decrease in the scale ratio enables the listener to evaluate more different individuals at the same time, and jump to a more distant place quickly. Conversely, increasing the scale ratio allows the listener to refine the existing musical pieces by evaluating more similar genotypes.

Shape change of the soundscape. A user can change the shape of the soundscape by modifying the genotype-phenotype mapping shown in Fig. 3. Every time the user shakes the iPhone or iPad, the position of each bit in the genotype that corresponds to each property of the sound sources is right-shifted cyclically by two bits. Then, the user jumps to the location on the soundscape based on the new genotype-phenotype mapping so that the sound sources of the genotype in the user's front left position are kept unchanged. This enables the user to explore a wide variety of mutants of the current population.

Bookmark of a location. If a listener touches one of the four buttons on the bottom left in the screen, they can save the current position, scale and direction of the soundscape as a kind of bookmark. One previously bookmarked state can be loaded as the current state of the user.

Finally, a listener can change some optional settings to facilitate exploring processes by touching the button on the bottom right.

2.3 Basic Evaluations

iSoundScape has been freely available from the AppStore[1], an online store for downloading applications for iPhone. Approximately 1,100 people around the world have downloaded it since May 2010. Users who have commented on iSoundScape have made several suggestions. The most important of these relate to orientation on the fitness landscape. The first suggestion is that stereo headphones or external speakers were effective for localizing sounds from right and left directions, although it was not always easy to localize sounds coming from in front of or behind the listener. Second, it was helpful to actively move back and forth on the soundscape for localizing the sounds. With these methods, listeners were able to effectively evaluate each musical piece and search for favorable musical pieces without difficulty.

The general consensus of the user comments was that the spatially grounded representation of musical pieces worked well on the soundscape and created a new kind of listening experience invoking a feeling that one is surrounded by different elements of a musical piece. A novel aspect of this experience is that a set of sound sources from neighboring genotypes is interpreted as a kind of musical work of itself. The wide variety of sound types including sound effects produced through the exploration process itself provides an entertaining and engaging experience because small changes in the sound types and locations can change the aesthetic feeling of the musical piece in unexpected ways. Also, the

[1] http://itunes.apple.com/us/app/isoundscape/id373769396?mt=8

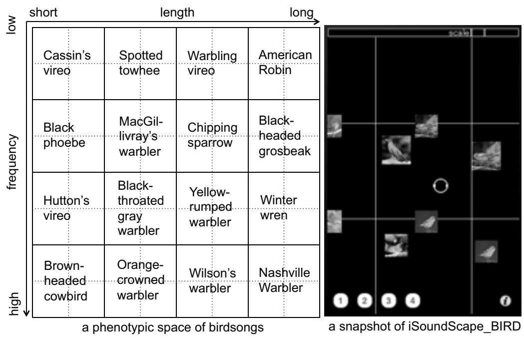

Fig. 4. A phenotypic space of birdsongs and a snapshot of iSoundScape_BIRD

scale or shape change of the soundscape enables a user to jump from the current location to distant locations at one time. This can be a helpful method to explore the whole soundscape quickly.

3 iSoundScape_BIRD

To take the system a step further, we developed a variant of iSoundScape termed iSoundScape_BIRD that features various birdsongs from California. There are several reasons to feature birdsongs. First, birdsong is one of the best examples of "nature's music", and may appeal to some people on this level. Second, there is the potential for practical applications of this system in exploring the ecology and evolution of birdsongs from a new perspective. For example, the parallel and simultaneous listening of songs from various directions is analogous to some aspects of natural soundscapes. Lastly, this system could provide a new way to teach concepts of adaptive evolution in an educational environment. Here, we enumerate features modified from the original iSoundScape.

We prepared 64 sound clips of birdsongs from the sound tracks recorded in Amador County, in Northern California. There are 16 species as shown in Fig. 4 (left), and 4 different song clips for each species. The duration of songs varied from about 0.5 to 7.0 seconds. The song properties varied significantly from this single location.

We have constructed two different soundscapes from these song clips. Each reflects a phenotypic or ecological property of birdsongs. In addition to providing an interesting way to explore birdsongs, the first soundscape was developed to provide a way to understand how properties of birdsongs could vary among species in a soundscape. We assumed a two-dimensional phenotypic space of birdsongs as shown in Fig. 4 (left). The axes in this space reflect roughly two basic properties of birdsongs: the length of songs and the intermediate frequency of songs. Instead of using genotypes, we directly mapped the all song clips to a 8×8 two-dimensional soundscape so that squared clusters of four song clips of each species are arranged according to the phenotypic space in Fig. 4 (left). Thus, each genotype is composed of a set of two integer values, each corresponding to the x or y location within the phenotypic space.

In a second form, we replaced the original sound clips in the iSoundScape with the clips of birdsongs. To do so, we assumed that the first 4 bits in the substring for each sound source in a genotype represent the species, reflecting the topology of the phenotypic space in Fig. 4, and the next 2 bits represent a unique number of song clips in the species. We also changed the number of sounds for each genotype from three to two. Thus, the soundscape comprises a $2^{16} \times 2^{16}$ grid. In this case, each genotype is composed of a 32 length bit string representing an ecological situation involving two birds, and is mapped to both individual birds singing different songs at different locations. Thus, a listener is able to explore the soundscape and virtual ecology of multiple individuals of songbirds.

In both cases, we inserted a random interval between the start times of songs from each bird to approximate the variation found in many natural soundscapes.

Finally, we used a small picture to represent the species of each individual[2]. This allows the listener to recognize the distribution of birds on the soundscape.

3.1 Preliminary Evaluations

Fig. 4 (right) shows a snapshot of iSoundScape_BIRD. In our preliminary evaluations with the phenotypic soundscape, we could recognize changes in the length and frequency of songs gradually through exploration of the soundscape, and understand how these species have different types of songs. In addition, it was helpful to recognize the difference between the similar but slightly different songs of different species because they are closely located on the soundscape and the listener can hear their songs from different directions at the same time to compare their properties between them. In the case of the ecological soundscape, the listener could feel the acoustic environment as more ecologically realistic because more birds were singing at different locations.

We believe both can provide a new way to understand ecology and evolution of birdsongs from scientific and educational viewpoints in addition to an artistic point of view. We are planning to release this variant at AppStore.

4 Conclusion

We proposed two variants of interactive evolutionary computation for musical works based on adaptive walk on a fitness soundscape. The first variant was developed to explore a new application of IECs for creating spatially grounded musical pieces as a new kind of interactive art, based on the concept of soundscape. Listeners were able to search for their favorite musical pieces by moving around

[2] These pictures are provided by Neil Losin and Greg Gillson from their collections of bird pictures (`http://www.neillosin.com/`, `http://www.pbase.com/gregbirder/`)

the soundscape actively, even within the limitations of two speakers for localization. The second variant was developed to explore soundscapes of songbirds to explore new ways to experiment with the ecology and evolution of birdsongs. It appears that this system may find application to better understand properties and ecology of birdsongs by creating phenotypic and ecological soundscapes.

Future work includes more detailed evaluations of these applications, a use of other sensors in the device such as GPS or an acceleration meter in order to make the system more interactive, and an addition of some functions that enable us to learn more about songbirds in the system.

Acknowledgements

The authors thank Zac Harlow, who gave fruitful comments and advice for development of iSoundScape_BIRD, and Neil Losin and Greg Gillson, who allowed us to use the beautiful pictures of songbirds from their collections. This work was supported in part by the Excellent Young Researcher Overseas Visit Program of the Japan Society for the Promotion of Science.

References

1. Biles, J.A.: GenJam: A Genetic Algorithms for Generating Jazz Solos. In: Proceedings of the 1994 International Computer Music Conference (1994)
2. Cherry, E.C.: Some Experiments on the Recognition of Speech with One and with Two Ears. Journal of the Acoustical Society of America 25, 975–979 (1953)
3. Dahlstedt, P.: Creating and Exploring Huge Parameter Spaces: Interactive Evolution as a Tool for Sound Generation. In: Proceedings of the International Computer Music Conference 2001, pp. 235–242 (2001)
4. Knees, P., Schedl, M., Pohle, T., Widmer, G.: Exploring Music Collections in Virtual Landscapes. IEEE Multimedia 14(3), 46–54 (2007)
5. Rocchesso, D., Bresin, R., Fernström, M.: Sounding Objects. IEEE Multimedia 10(2), 42–52 (2003)
6. Sims, K.: Interactive Evolution of Dynamical Systems. In: Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life, pp. 171–178 (1992)
7. Suzuki, R., Arita, T.: Adaptive Walk on Fitness Soundscape. In: Proceedings of the Tenth European Conferences on Artificial Life (ECAL 2009). LNCS, LNAI (2009) (in press)
8. Takagi, H.: Interactive Evolutionary Computation: Fusion of the Capabilities of EC Optimization and Human Evaluation. Proceedings of the IEEE 89(9), 1275–1296 (2001)
9. Unemi, T.: SBEAT3: A Tool for Multi-part Music Composition by Simulated Breeding. In: Proceedings of the Eighth International Conference on Artificial Life, pp. 410–413 (2002)
10. Wright, S.: The Roles of Mutation, Inbreeding, Crossbreeding, and Selection in Evolution. In: Proceedings of the Sixth International Congress on Genetics, pp. 355–366 (1932)