STUDENT ARTICLE

An introduction to Discovery Strategy

AMBROSE FIELD

Department of Music, City University, Northampton Square, London EC1V 0HB, England

The purpose of Discovery Strategy is to define compositional methods which can lead first-time listeners to the musical substance of an electroacoustic work. Discovery Strategy provides a steering mechanism which directs first-time listeners towards identifying features that lie beneath the surface of the music. A key component of this mechanism is the skilful integration of familiar sounds into more abstract contexts. The implementation of the strategy rests on a consideration of how different listeners perceive music through time, and how their perception changes over repeated listenings. The tape work *Undercurrents* (Field 1994) follows these principles, demonstrating that it is not necessary to compromise structural sophistication while pursuing Discovery Strategy methods.

1. LISTENING

The members of an audience at a public electroacoustic concert will usually range from those who have never heard electroacoustic music before, to practised analysts and composers. Different listeners will have varying degrees of knowledge about the structures of electroacoustic music, and those with greater knowledge will be able to recognise correspondingly more subtle elements of the music. Truax (1984: 48) introduced the notion of musical competence as 'the tacit knowledge a musician has about musical structure'. This notion can be extended to the listener and will be referred to as 'listening competence'.

Having recognised that different listeners possess differing levels of competence, it is necessary to analyse ways in which people perceive sound. Schachtel (in Smalley 1992) identifies two main modes of auditory perception – autocentric and allocentric. In brief, autocentric listening focuses on emotional responses to a soundscape, including the feelings of pleasure and displeasure. However, during allocentric listening the listener gains enjoyment directly from the structures and features of the prevailing soundscape, regardless of any emotional response they may provoke. Although all listeners are capable of both modes of listening, the degree of allocentricity is governed by the subject's listening competence.

Smalley (1992) notes that 'Autocentric attitudes to music and sound are common'. If this form of listening is so common, one should question why it is so often ignored by composers of innovative contemporary music. The answer may lie in the composer's ability to predict audience reactions based on each mode of listening. A listener's allocentric reactions may be quantified in terms of strictly musical parameters. However, due to their primitive, emotive nature, autocentric responses are perceived as being more difficult to control or altogether irrelevant. Electroacoustic music presents a unique opportunity for composers attempting to systematise the listener's autocentric reactions. This is possible because electroacoustic compositions may include sound materials that provoke direct emotional responses by means of their extra-musical implications (e.g. a fire alarm).

Autocentric and allocentric listening modes are not mutually exclusive. By exploiting an initial autocentric response to familiar sound materials, the composer can steer listeners towards an allocentric perception of features deeper in the structure of the music.

Before discussing the steering mechanisms themselves, it is necessary to examine a single prerequisite: it is essential for a work employing Discovery Strategy methods to have a highly unified, multi-levelled structure.

1.1. Necessity for a multi-levelled structural approach

For analytical purposes, it is often necessary to visualise music as consisting of many levels. The concept of levels is linked to the experiential timescale of a work. If a certain trend or feature can be perceived over a short time span, that feature is said to be on the surface level. A longer time span is required to ascertain whether the recognised features can be grouped into more general trends. It may require many listenings to establish which features exist on a background level. Figure 1 shows how timescale affects feature recognition within *Undercurrents*.

Discovery Strategy places an important condition on the use of multi-levelled structures. Since listeners

Organised Sound **1**(2): 119–23 © 1996 Cambridge University Press

https://doi.org/10.1017/S135577189600009X Published online by Cambridge University Press

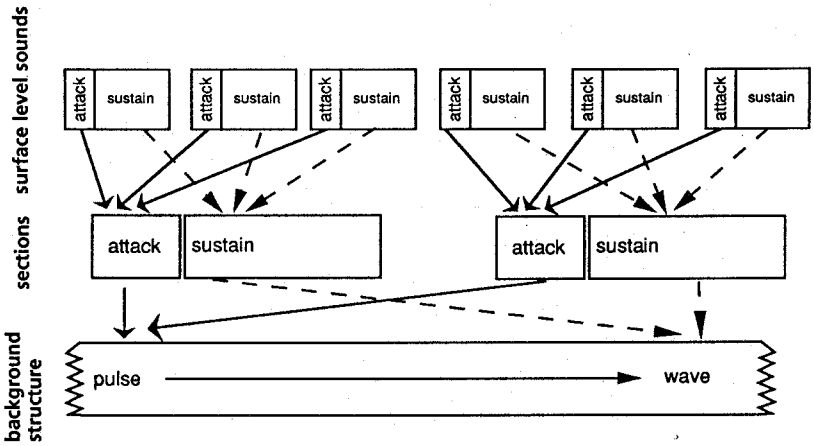


Figure 1. The sounds have two components: a short pulse, and a protracted decay. They are used within sections which have two similar components: a short group of closely spaced attack-morphologies, followed by a dense protracted decay.

have to be encouraged to return to the music, and since the method of doing this is by making a direct link across structural levels, it is necessary that the structural poietics of each level are similar.

Taking *Undercurrents* as an example, figure 2 shows how the structural principle of change from pulses to waves operates on every level of the piece. A direct link is established between the surface transformation and the overall background structural logic. This process is defined as cross-hierarchical linkage. In *Undercurrents* the link is accomplished by a thinning-out of the overall texture to focus the listener's attention on the transformation itself (audio example 1). Today's computer-based audio mixing environments allow this process to be carried out with a degree of control hitherto not possible.

The advantage of using one central structural process for a work is that the perceived background structure varies little, regardless of the competence of the listener or the number of times the music is heard. Once such a situation can be established, the composer can attempt to steer the listener's attention towards selected features of the music that would otherwise require subsequent analysis.

Discovery Strategy also dispenses with the concept of a musical work as a single expanse of time, with structures imposed by the composer. Instead the work is visualised as many possible expanses of time, the structures within each of which depend on the interpretation of the listener. It is the job of the composer to help listeners to interpret the work, aiding them to discover the most pertinent features at the right times.

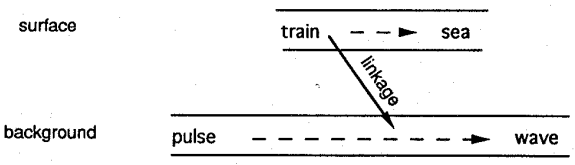


Figure 2. On the surface level, a pulsed sound transforms into a wave-shaped sound. During this transformation, other associated materials melt away, promoting a background level structural process to the foreground of the listener's attention. If the structural processes on each level were not the same, it would not be possible to make this linkage.

2. STEERING PROCESSES

A steering process aims to direct a listener's attention towards, or away from specific features of the music. Steering processes fall into two groups, according to the experiential timescale on which they operate.

Primary steering processes function mainly on a first audition, and encourage allocentric listening. This is achieved by provoking an initial emotional, autocentric response, through the use of sound materials with which the listener can identify.

Secondary steering processes are those that direct the listener's attention towards deeper structural concerns, simultaneously providing a mechanism for the composer to control the manner in which different features of a work are revealed over repeated listenings.

2.1. Primary steering processes: the integration of environmental sounds into abstract contexts

Environmental sound is a powerful phenomenon. It can be recognised as being part of human experience, regardless of the level of listening competence. If the ear can be led skilfully from recognisable environments to those which are aurally unfamiliar, an opportunity exists to give first-time listeners an easily accessible route into the piece without compromising the enjoyment of the music for more practised listeners.

Before proceeding it is necessary to define precisely what is meant by environmental within the context of this discussion:

https://doi.org/10.1017/S135577189600009X Published online by Cambridge University Press

> Environmental refers to recorded sounds which are not processed in any form that would prevent a listener from identifying the recording as matching an equivalent sonic reality. The anecdote, or referential qualities of the sound are strongly preserved. With sounds of this type, listeners first focus on identifying the source and cause, simultaneously building a personal meaning from what the sound signifies. These sounds may also be referred to as anecdotal sounds. (Field 1995)

To implement this steering technique, two well-known electroacoustic sound treatments can be used.

2.1.1. Interpolative linear transformation

An interpolative linear transformation is a smooth transition from one feature to another over time. As the name suggests, the gradual transition that occurs in the interim is generated by a process of spectral interpolation. As a sound treatment, the process is most clearly highlighted by Trevor Wishart's program VOCINTE, for the Composer's Desktop Project.[1] Although this is a very useful tool, a number of problems arise which must be considered before constructing a steering process.

(1) Regardless of the material, the perceived result of an unsuccessful interpolative transformation is a passage from sound one, which at some point in time flips directly to sound two. Perversely this annuls the effect the treatment was intended to create. Such a situation would unhinge environmental materials from their supporting abstract structures by directing the listener's attention towards the medium rather than the message.
(2) As the process is essentially a progressive spectral exchange,[2] the materials subjected to transformation must be closely matched in the frequency domain for the transformation to succeed at all. For example, this would discount the possibility for transformation between fine, low-density granular textures and a live recording of high frequency bird song. When linear transformation is attempted under these conditions, the result sounds little different from a simple cross-fade between the materials.

Clearly, it is not possible to depend on interpolative transformation alone to effect the steering process. Instead, it is useful to consider the mapping of the gestural characteristics of the environmental material onto more abstract sounds.

2.1.2. Gesture linkage

As musical gestures often result in distinct patterns of pitch and dynamic motion through time, gesture linkage encompasses the mapping of the pitch and envelope contours of one sound onto another without an exchange of spectral information.[3] If the shape of an abstract sound can be made to be reminiscent of one that is environmental, the possibility arises for a structural link between the two.

This technique is used only once in *Undercurrents:* towards the start of the closing section the sound of the sea is given the pitch and articulation of an earlier train sound. However, this instance also demonstrates the instability of this technique. Due to the pitch shift, the result of the treatment sounds much like a car passing, making it difficult for the listener to hear this feature as a hybrid combination. This is because synthesised or artificial sounds can often sound more real than the actuality.

Although the process of gesture linkage is readily applicable to time domain sound transformations, a more subtle application is to map low-level structural trends onto specific features of the sonic surface. This technique is of fundamental importance within Discovery Strategy. The mapping of long-term structures and surface morphological design within *Undercurrents* is shown in figure 3.

The composer can steer the listener away from anecdotal, autocentric listening by setting up an appropriate spectral context *before* the introduction of the environmental sound. Two rules appear to govern this procedure:

(1) Sharply contrasting materials of differing environmental/abstract types will fail to be integrated successfully. In general, it is beneficial if the materials are suitable for the linear transformation process mentioned previously.
(2) The structural logic of preceding materials must be compatible with the spectral logic and referential qualities of the environmental sound being introduced. If the anecdotal sound can exhibit the same structural logic when part of a new context it can be said to be transcontextual. The new context itself must have a similar spectral logic to the one from which the anecdotal sound has been abstracted. For example, if a sound is part of a

[1]VOCINTE (CDP 1994) is a spectral interpolation program which operates on the results of phasevocoder analyses of two sounds. Over a period of time specified by the user, the amplitude and phase components of two sounds are separately interpolated. IRCAM's SVP1.1 can be used to similar effect using the 'G-cross' cross-synthesis mode.

[2]New work by Tom Erbe and Larry Polansky (Erbe and Polansky 1994) has attempted to build mutation trajectories for sounds by spectrally bending the source sounds towards a target sound. This process is best exemplified by Tom Erbe's program Soundhack version 0.8.

[3]Alchemy, by Passport Systems, is perhaps the most useful commercial software with which to accomplish gesture mapping due to the provision of graphical tools for pitch and envelope extraction.

https://doi.org/10.1017/S135577189600009X Published online by Cambridge University Press

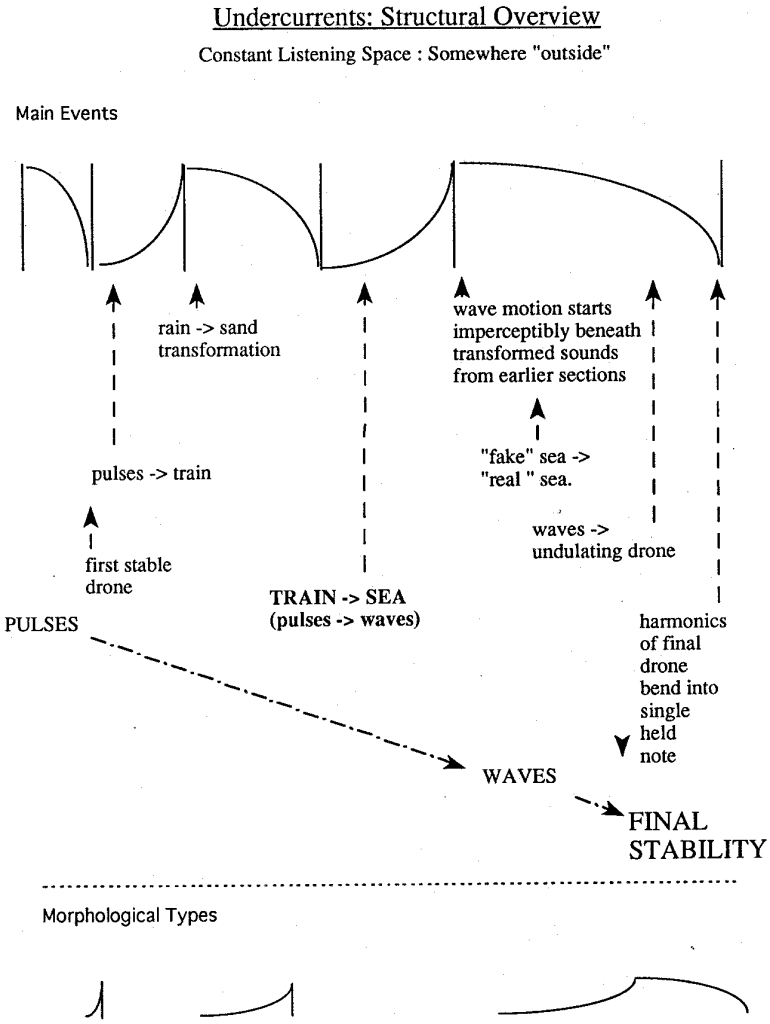


Figure 3

sonic agglomeration process in the original instance, it could be integrated within an abstract environment where the same process is already active.

Perhaps Georges Braque was aware of this kind of transcontextual logic when painting *Atelier VIII* (1952):

> No object can be tied down to any one sort of reality; a stone may be part of a wall, a piece of sculpture, a lethal weapon, a pebble on a beach, or anything else you like, just as this file in my hand can be metamorphosed into a shoe-horn, or a spoon, according to the way I use it. (Braque, as quoted in Tucker (1992))

Braque is not talking about simple context montage: the stone remains the stone in a variety of situations, but unlike many montage contexts, it does not look out of place within a new environment.

2.2. Secondary steering processes

Secondary steering processes help a listener to maintain an allocentric listening attitude. They also attempt to allow selected elements of a composition to be exposed over successive listenings. This can be accomplished by building on two perceptual premises. They are:

(1) When hearing a new feature listeners process the most readily available information first. This allows one to gain useful knowledge about a situation very quickly, and is often the reason why surface level features are the first to be detected.
(2) Listeners internalise experience in the most economical way. For example, if the first minute of a piece consisted of a series of sharp attacks, most listeners will remember that a pattern of similar features occurred, rather than remembering every instance of an attack separately.

When we listen to music our perception is learning. It is learning what are the most identifiable features of the musical surface. When we listen again, we already know what the main features are. This assimilation will undoubtedly be different for different people, and will vary depending on their listening competence. To understand how a composer can

https://doi.org/10.1017/S135577189600009X Published online by Cambridge University Press

actively influence the perception of musical structure when a work is repeated, it is necessary to analyse some auditory processes that can be exploited musically.

2.2.1. Redundancy

In information theory, the multiple representations of the same feature within a specific local context generate redundancy. Compositionally, this can apply only to features within a single structural level. If there is a group of identical materials, they can be represented by the same symbol when placed at a deeper structural level. It is possible that the concept of redundancy is highly relevant to the human decoding of the surface of music. We do not have to apply any more energy to decode and store features that we have already learned – we merely know that they still exist.

Redundancy is a key device explored in the opening moments of *Undercurrents* (audio example 2). The first attack surprises most listeners, and it is reiterated three times. The level of small, local detail between these repetitions is progressively increased. The listener should be able to hear more detail as the attention-commanding power of the attack diminishes due to the increasing perceptual redundancy. It seems that only a certain level of redundancy can be tolerated until a composition becomes dull.

2.2.2. Gestalt grouping of feature types

This process is very similar to redundancy, but does not involve identical materials; it involves the grouping of similar materials. It seems that this would require more mental effort to detect than redundancy, but in fact it does not. In the process of detecting whether one material is exactly the same as another, it is the generality that is matched first. On the first listening, the power of the opening of *Undercurrents* comes from a pattern of sharp attacks followed by a more constant section. Similar sounding materials after the attack may be perceptually grouped as a single unit. Together these events form a large-scale onset/continuation morphology, as shown in figure 1.

According to acoustic theory (Winkler 1967), we gather most timbral detail in the onset stage of a sound. However, as the textures following the attacks could be grouped easily by general type, they have been laden deliberately with small, specific spectromorphological detail. It is anticipated that most listeners will group these sections together on a first listening. Only on a repeated hearing will the textural detail start to emerge.

It is these two processes of redundancy and gestalt grouping which form the core of the secondary steering tools within Discovery Strategy.

3. CONCLUSION

The major components of Discovery Strategy have been summarised. The methods introduced allow composers to steer the listener's attention through a work. These systems operate actively over many repeated hearings. Importantly, Discovery Strategy does not alienate those who have an advanced knowledge of the sounds and structures of electroacoustic music.

REFERENCES

Composers' Desktop Project (CDP). 1994. *Professional Computer Music System User Guide*, Release 2. York: CDP Ltd.

Erbe, T., and Polansky, L. 1994. *Soundhack .80 User's Manual*. Lebanon, USA: Frog Peak Music.

Field, A. 1995. Electroacoustic composition. *Unpublished PhD Dissertation*. City University, London.

Smalley, D. 1992. The listening imagination: Listening in the electroacoustic era. In J. Paynter, T. Howell, R. Orton and P. Seymour (eds.) *Companion to Contemporary Musical Thought*, Vol. 1, pp. 514–54. London and New York: Routledge.

Truax, B. 1984. *Acoustic Communication*. Norwood, New Jersey: Ablex Publishing Corporation.

Tucker, M. 1992. *Dreaming with Open Eyes*, 38. Musselburgh, UK: Scottprint.

Winckel, F. 1967. *Music, Sound, and Sensation*, 170. New York: Dover Publications.

https://doi.org/10.1017/S135577189600009X Published online by Cambridge University Press