
Space in sound: sound of space

DAVID WORRALL

Australian Centre for the Arts and Technology, Australian National University, Canberra 2600, Australia
E-mail: David.Worrall@anu.edu.au

I think that all music will become space music and that space becomes as important as pitch in the traditional music, as durations and rhythm and metre and there is a very new development of harmony of space and I mean space chords, space melodies and that doesn't mean pitches, it means movement on several levels around the listener: above, below, in all directions. (Stockhausen 1997)

1. INTRODUCTION

At a time when our sense of physical space is being radically challenged and modified by new global communication technologies, it seems not unusual that composers are taking a renewed interest in it. We have begun to experiment with using the new technologies available to us to move sound in space and create surround-sound environments.

For electroacoustic music and other soundscapes we have pan pots on our mixing desks and in our synthesizers to radially position sounds between loudspeakers, and reverberators to simulate the distance of sound sources from the listener. Dolby Digital 5.1 (AC-3) and DTS were introduced as competing standards for theatre sound recordings around 1992 and consumer products are available in both formats. In addition, Digital Versatile Discs (DVDs) are making sophisticated soundfield techniques such as Ambisonics both practical and more accessible.

This article traces some of my own explorations in the use of 3-space¹ for musical composition, examines the limitations of basing structuring methodologies primarily on functional psychoacoustic studies of hearing, and suggests alternative approaches based on an understanding of 3-space from the work of the perceptual psychologist James Gibson.

2. PERSONAL EXPERIENCE

My own interest in 3-space as an aspect of musical composition began in 1981 with *Mixtures and Recollections*, my first composition by direct computer synthesis. *Mixtures* is a quadraphonic composition

¹ Concrete space, or real space (hereafter called 3-space) defined by the *OED* as: '1. Continuous extension viewed with or without reference to the existence of objects within it; 2. Interval between points or objects viewed as having one, two or three dimensions.'

which relies on the fact that 3-space acts as a strong stream segregator; sounds located by a loudspeaker at the same physical place at the same time are associated (see Bregman 1990: 75–6). Three layers, transposed in pitch and time compressed, revolve around the auditorium at different velocities in a kind of spatial canon. What one hears is the spatial accretion of the layers into multiple musical gestures in different locations in the 3-space. A stereo mix of the work even further de-emphasises the original layering.

Later at ACAT we developed a 16-channel composition/performance environment (Worrall 1989) which used a 'brute-force' approach to moving sounds in 3-space for works such as *Life Dreaming* and *Cords*. Sounds for each of the loudspeakers, equally distributed over the surface of a geodesic hemisphere, are generated independently under MIDI control. These experiments have led to the construction of an ambisonic distribution system (Vennonen 1994) which is still being further developed (Vennonen 1995) and refined to include realtime encoding and decoding of multiple audio signals.

3. DEVELOPING A RELATIONSHIP BETWEEN 3-SPACE AND OTHER CHARACTERISTICS OF SOUND

A prime concern during the development of this performance environment was how to integrate 3-space into different composing methodologies. I began, perhaps naturally enough, to think of 3-space as a locational aspect of sound which could be added as another dimension to an abstract² multidimensional musical space of pitch, duration, dynamic and timbre:

There remains a fifth dimension, which is not, strictly speaking, an intrinsic function of the sound phenomenon, but rather its index of distribution: I refer to

² Abstract space, defined by the *Penguin Dictionary of Mathematics* as: 'A formal system characterised by a set of entities, together with a set of axioms for operations on and relationships between these entities (e.g. metric spaces, topological spaces, and vector spaces).'

space. Unfortunately it was almost always reduced to altogether anecdotal or decorative proportions, which have largely falsified its use and distorted its true functions. (Boulez 1971: 66)

Boulez, speaking about ensemble writing, says:

... spatial distribution seems to me to merit a type of writing just as refined as the other sorts of distribution already encountered. It ought not only to distribute spaced-out ensembles according to simple geometric figures, ... it must order the micro-structure of these ensembles.

It is obvious that the index of distribution, space, acts not only on the durations, but also on pitch, dynamics and timbre within the time-span; static or mobile distribution can be considered as maintaining certain relationships with all these interacting characteristics. These relationships are of greater subtlety than those of simple speeds – angular or lateral according to the spatial layout in use, even leaving out of account the *local* acoustical conditions, which seriously complicate the problem. ... It seems to me that the real interest in distribution lies in the creation of ‘Brownian movements’ within a mass, or volume of sound, so to speak; hence it is a question of elaborating a strongly differentiated typology of relationships, to be set up between the phenomenon itself, whether individual or collective, and its actual, absolute place in real space. (Boulez 1971: 66–7)

This use of 3-space is similar to the way that timbre was used to carry pitch–time constructions in the Baroque and, as Trevor Wishart has eloquently elaborated (Wishart 1996), is fraught with difficulties, not the least of which is the way it imposes undesirable limitations, mainly derived from the limitations of notation, especially in making works using techniques for transforming the sonic material directly. Just as it has become compositionally useful to find more interesting ways to explore the nature of timbre space, it has become important to find more appropriate ways to explore the nature and use of 3-space so as to avoid using it purely decoratively in effects such as sounds ‘whizzing’ around the auditorium.

4. TIMBRE SPACE AND 3-SPACE: SIMILARITIES AND INTERSECTIONS

Both timbre and 3-space are multidimensional and are psychoacoustically related. Timbre is a difficult term and is often defined by negation, as Begault suggests:

The spectral content (spectrum) of a sound source, along with the manner that the content changes over time, is largely responsible for the perceptual quality of **timbre**. Sometimes referred to simply as ‘tone color’, timbre is actually difficult to define. A ‘negative’ definition that is commonly cited for timbre is “the quality of sound that distinguishes it from other sounds of the same pitch and loudness”. (Begault 1994: 35)

Interestingly, he goes on to suggest:

This could be extended to spatial hearing, in that two sounds with the same pitch, loudness, and timbre will sound different by virtue of placement at different spatial locations. (Begault 1994: 35)

Both timbre and spatial cues depend on the morphology of the spectral modification (emphasising some frequency bands whilst attenuating others) of the sound source. Theile (1986) and others have suggested that it is not the stimulation of the eardrum that determines timbre, but rather the overall sense of hearing, a higher-level brain function that identifies the timbre and location of the sound source. Clearly the two interact; modifications of timbre (including all its sub-dependencies such as fundamental frequencies, amplitudes, spectra, etc.) affect our perception of different aspects of location: proximity (distance), angle (lateral (left–right, front–back)) and azimuth (height) and, for moving sources, velocity.

In his overview article on music representation systems, Roger Dannenberg suggests that 3-space is an aspect of timbre:

With timbre we are still learning what to represent. My explanation is that, starting with sound, we picked out the two things we understood, pitch and amplitude, and called everything else timbre. So timbre is by definition that which we cannot explain. As aspects of timbre are isolated and understood, such as spatial location and reverberation, these components come to be regarded separately, leaving timbre as impenetrable as ever.

Taking a less cynical view, real progress has been made towards timbre representation. The classic studies by David Wessel (1979) and John Grey (1975) used multidimensional scaling and refined the notation of the timbre space. Although these studies represented timbre in terms of perceptual dimensions, others have represented timbre in terms of control dimensions, such as the set of parameters needed for a particular synthesis algorithm.

The current practice is usually to represent timbre by name and number. Real timbre space has so many dimensions that it is often preferable to locate and name interesting points in the space. MIDI program change messages and the concept of ‘instrument’ found in many software synthesis systems are examples of this approach. (Dannenberg 1993: 23)

In keeping with Wishart’s warnings on the way notational concerns impact on our perception of the importance of different aspects of the multidimensionality of sound, we need to be cautious in suggesting any hierarchical relationships between 3-space and timbre:

It is notability which determines the importance of pitch, rhythm and duration and not vice versa ... much can be learned by looking at musical cultures without a system of notation. (Wishart 1996: 6–7)

In fact timbral modifications often suggest spatial modulations and this could suggest that spatial perception happens even later in perceptual processing than timbre, as it radically depends upon it. At best we can say, like most abstracted parameters of multi-dimensional musical spaces, they are inextricably linked.

Bearing in mind Wishart's appropriate cautioning about timbre:

A catch-all term for all those aspects of a sound not included in pitch and duration. Of no value to the sound composer! (Wishart 1994: 135)

the instrument concept is useful because it affords 'toolness', and in our own current development of an ambisonic performance environment, referred to earlier, we are exploring the MIDI sampler model as a way of representing 3-spatial movement by continuous controller patterns which are stored (in the ambisonic spatialiser) as patches in breakpoint tables (time-offset, lateral angle, azimuth, proximity, velocity). These tables can then be interpolated, compressed or expanded, permitting spatial patches with global time independence. Many difficulties arise, however, in trying to systematically translate the interaction of these parameters with other components of a musical gesture into variegated yet coherent 3-spaces, suggesting that more sophisticated tools are needed.

5. HEARING COMPARED WITH LISTENING

Most models of aural perception, firmly rooted in Cartesian dualism, rely on a functional analysis of the physiology and psychophysiology of hearing. These studies, such as those by Fletcher and Munson who mapped curves of equal loudness, are relevant to the sound designer, especially for tool building, and play an important role in their contribution to an understanding of the mechanisms of hearing.

Functional studies of the ability of humans to locate sounds in 3-space, mostly using broadband noises and sine tones with the test subject unable to move their head, are used to describe what are known as binaural head-related transfer functions (HRTFs). HRTFs can be thought of as maps of frequency-dependent amplitude and time differences that result primarily from the interaction between the head and the complex shaping of the ears' pinnae. Each individual has a unique binaural HRTF although there are some generalisable characteristics. Begault's recent book (Begault 1994) contains a thorough summary of this research which has received added impetus from the growing interest in the role of audio in multimedia and virtual reality, as its title attests.

Even though each ear receives different auditory signals, the auditory 'image' is one of objects that

appear coherent in the outside world and not in the ears or head. This is also true despite head and body movements, so the observer must be dynamically compensating for the position and orientation of the head, and of course it is at least possible that the muscles of the neck and the different reflective properties of the upper torso also play a part. It is perhaps useful to remind ourselves that we evolved as integrated organisms and, despite Platonic idealism, our heads are only one part of our bodies!

Given the role of post-mechanical (neural and mental) analysis in timbre and space perception, the *what* must at least be as important as the *how*. Whilst psychophysical studies are useful in describing how one hears these types of test noises in very controlled environments, they are of limited value outside of them. More particularly, they tell us nothing of *what we are listening*. Stephen Handel puts it like this:

Listening is not the same as hearing. The physical pressure wave enables perception but does not force it. Listening is active; it allows age, experience, expectation, and expertise to influence perception. It is often helpful to illustrate how the ear is like a microphone or how the eye is like a camera. It is a mistake, however, to equate the ear with listening or the eye with looking, or to equate the faithful recording of sound energy or light energy with hearing or seeing. We hear and see things and events that are important to us as individuals, not sound waves or light rays. . . . The study of listening must take place within the context of the environment in which listening *evolved*, since it is the product and reflection of that environment. (Handel 1989: 3)

It is clear, then, that a general theory of spatial cognition that can be used to integrate data from diverse disciplines must address the functions of all relevant perceptual and cognitive structures and processes, and we must be careful not to automatically eliminate any of the parts purely on the basis of a philosophical attitude.

6. THE JASTREBOFF MODEL OF HEARING³

The Jastreboff model of hearing, which follows the 'digestive' model of perception as summarised by Kolars (1972: 192), moves away from the electro-mechanical model of the ear – especially of the function of the auditory nerve. The mechanical ear changes sound waves into electrical patterns which are passed along the auditory nerve to the part of the cortex in the temporal lobe of the brain. This is quite a long way from the ear and there is no perception of sound until these electrical impulses reach it.

The auditory nerve is a bundle of approximately 10,000 fibres made up of millions of cells. Between

³ This section was developed from a radio interview with Jonathan Hazel, The Royal Institute for Deaf People, London, 'Ockham's Razor', ABC Radio National, August 1997.

the ear and the temporal lobe there are subconscious pathways which consist of millions of nerve cells joined together by many connections in dense networks that act like filters for identifying patterns of sound. These networks group together frequencies on the basis of 'knowledge' of how sounds form complex signals. The subconscious part of the brain has been taught to recognise and respond very strongly to important sounds (such as one's own name, and timbral forms) by 'encoding' an amplifier for them on the auditory nerve so that we can hear them more clearly and categorically, and thus respond to them more quickly.

So, according to this model, as the signal moves along the auditory nerve, it is being subtly but significantly modified by all sorts of influences that are crowding around this signal pathway as it passes towards the brain. There has to be a lot of processing of this information because the auditory nerve just passes along frequency information. It does not know what speech is, what environmental sound is, what internal body sounds are. So these subcortical pathways have to sort out and group those frequencies and identify what the discrete signals are. This is a very complex processing procedure but once it is completed, those signals which are meaningful are enhanced and those signals which are not are reduced in amplitude or even filtered out.

These conditioned reflexes are initiated by emotional responses from the limbic and autonomic systems of the brain which cause a heightened awareness. This heightened awareness creates a survival reflex and autonomic response (very like the fight or flight response to fear, for example) and cements this group of frequencies in the auditory pathways in the way suggested. This autonomic response is an explanation for why we can easily recognise timbres, once 'imprinted', so directly.

Spatial characteristics are more complex in that they also involve more disparate and environmentally derived parameters such as our expectation that low frequency sounds are likely to be closer to the ground. Yet they are hardly less important from an evolutionary viewpoint. There is not much point in being able to recognise the call of a predator if you cannot tell from which direction and how quickly it is coming! However, it is likely that processes similar to those described for timbre recognition are in operation.

7. THE PHENOMENAL OR NATURALISTIC PERCEPTION OF 3-SPACE

The evidence used by phenomenal or naturalistic approaches to perception consists of one's conscious experiences. 'Naturalistic' means that the evidence

concerns responses to whatever stimuli occur naturally within the environment; there is no attempt to modify these stimuli or create artificial ones. Probably the most radical and to a certain extent least formal work in this vein is that of James Gibson. Developed from Aristotelian empiricism, his approach (Gibson 1966) deals in considerable detail with certain aspects of 3-spatial orientation such as depth perception and what he calls an 'ecological approach', which focuses on the process of adaptation between organism and environment. Thomas Lombardo explores the evolution of this idea of reciprocity in detail. For Gibson,

The structures and capacities of animals were described relative to their ways of life within an environment; in turn, the environment was described relative to the ways of life of animals. An explanation of perception involved a dynamic interdependency of animal and environment. Gibson's epistemology is direct realism, as was Aristotle's – the "object" of perception is the real world, viz., the environment. Perhaps the most difficult and unique point to grasp is treating perception as an ecological phenomenon rather than a mental or physiological event, yet Gibson's direct realism only follows if perception is defined ecologically. Perception does not reside in the brain or the mind any more than life resides in cells or in some inexplicable living spirit. Neither mentalism nor physicalism is correct. Perception, as well as life, is ecological; perception exists at the reciprocal interface of animal and environment within an ecosystem. (Lombardo 1987: 5–6)

Gibson's early work was in the study of depth perception and is critically reviewed by Bruce Goldstein (Goldstein 1984). His ideas were developed with particular reference to visual perception and a thorough and sensitive translation of them into the aural domain is needed. This early work is based on three main ideas; the ground theory of space perception, invariants, and direct perception, and is expanded into a mature theory of perception in Gibson (1979).

7.1. Ground theory

According to the ground theory, information contained in the ground (usually horizontal) plane is a texture gradient. The elements that make up a textured surface appear to be packed closer and closer together as the surface stretches into the distance; there is more texture detail the closer the object is to the observer. This gradient results in an impression of depth, and the spacing of the gradient's elements provides information about the distance at any point on the gradient.

For sound, Gibson's ground roughly equates with background ambience and texture roughly equates with reverberance, which causes the texture of a sound to be more indistinct the further away from

the auditor it is. Along with reverberation, texture gradients share other depth cues such as relative loudness (more distant elements of the gradients get softer) and spectral profile. A texture gradient extends over a large volume of 3-space and no matter where one moves on the gradient, the elements of the gradient provide information that enables one to determine the distance between different locations anywhere else on the gradient. This is in keeping with our known ability to perceive the difference between a soft sound emitted close to the listener and a loud sound emitted at a distance, even when they have the same amplitude at the ear.

One consequence of this theory is that if we present sounds in different virtual reverberant spaces closely in time, our ability to locate their virtual proximity could be considerably impaired. However, if different sounds are presented within a single virtual reverberant space, they will be more likely to be perceived as different sounds because they will have different textures consistent with the overall texture gradient of the space. This accords with the findings of John Chowning (Chowning 1971) who found the use of two types of reverberation, both local and global, resulted in superior location detection of computer-generated sounds. In addition, the more complex and variegated the background ambience is, the more likely one will be able to locate nonambient sounds because these nonambient sounds will be heard relative to the ambience: the ambience creates the space in which other sounds are heard.

7.2 Invariants

When an observer moves relative to a texture gradient, the texture will be in constant flux to the ears, i.e. the contours that define the textures of the gradient change, but some information – the texture of the gradient – remains, in Gibson's terminology, invariant and thus the scale of depth in the scene remains constant. According to Gibson, it is this invariance information that we use in everyday perception as we move through the environment.

Another example of invariance is provided by the way movement of an observer causes the textures of sounds in the environment to flow. This flow of the environment illustrates how, when the environment speeds past an observer who is travelling, there is flow of other sounds in the environment which is everywhere different to the way it flows from the location towards which the observer is moving. Since this point is at the centre of the flow pattern, it does not change (at least as much), i.e. it stays relatively invariant. Aurally, this change in relative invariance is caused by the front-back asymmetry of the body (including but not exclusive to the head) as mapped

in an HRTF, by Doppler shift and by the comb-filtering effects caused by the different reflective properties of the variegated environment. Thus, to 'stay on course' while moving toward a sonic source, one only needs to keep the invariant centre of the flow pattern centred on the source.

A consequence of this theory is that it may be possible to create the illusion of a sound moving towards a listener by keeping the sound-ground invariant and changing the texture of the sound. To create the illusion of the listener moving towards a sound, a change in the texture of the ground is needed as well. I have yet to test this postulate – perhaps it is possible to create the aural illusion of moving backwards or sideways or even perhaps to dance!

7.3. Direct perception

Perhaps the most controversial of Gibson's ideas is his explanation of how an observer uses this invariant information. His answer is that we pick up invariants by direct perception and just use them. We are pretty good at perceiving the correct sizes of objects in the natural environment, even if these objects are located at different distances from us. Gibson feels that the information provided by invariants is present in a form such that an object's size in our field of view *and* its distance from us can be perceived directly and immediately from an unconscious mental process of estimating how many units of a texture gradient are covered by the object without needing to be processed in any way. Gibson's explanation works well when we can observe both object and background. However, some researchers have noted that he cannot explain how we can perceive movement when the background is not visible, as in the case when a spot of light is seen in the dark or when an object is seen against a completely textureless background. In these cases, the background cannot provide the information we need to decide whether the object is moving or we are moving.

8. SPACE IN SOUND

3-space is not an unvariegated abstraction, it is defined by coherent combinations of objects and surfaces and what they afford us. Sound is not an abstract ideal which is projected into 3-space. The space is *in* the sound. The sound is *of* the space. It is a space of sound, but there is no sound *in* space. With experience we develop a memory of invariant characteristics that enables us to mentally separate the source (the invariant qualities) from the environments in which we hear the source. Of course we do this all the time when listening to recordings of sonic spaces: these recordings are overlaid or interlaced

with an existing sonic ambience which we can often 'mentally' filter out without too much difficulty.

The functional relationship between the mind and the body in this scenario is important because it is, I suggest, a vital part of the bridge between hearing and listening. It seems unlikely that the epistemology of indirect perception which follows historically from Plato's dualistic ontology⁴ can provide an adequate explanation for this. However, Gibson's theories of direct perception account for the reciprocity of a perceiver and their environment in a unified perceptual system. These theories are Aristotelian in concept and are supported in audition by the Jastreboff model of hearing which explains the way the auditory system 'learns' and adapts to the sonic environment in which the perceiver exists.

Naturalistic theories of perceptual psychology, based on observation in complex environments, offer significant insights into our understanding of 3-space not offered by other methods, and these insights, combined with stream analysis, suggest ways of organising and modifying sounds to enhance and integrate 3-space into compositional practice at a more fundamental level.

9. HEADPHONES, LOUDSPEAKERS, REVERBERATION AND CYBERSPACE

There is a major perceptual difference between sounds presented to the human auditory system over headphones and sounds presented in 3-space. In the former case, the experience is of the sound being inside one's head, yet this concavity is not experienced when using loudspeakers where the listener is free to move independent of the sound sources against a coherent ambience (background).

Perhaps it is precisely because of the variegated yet coherent nature of the sonic ambience of real 3-space (it has no internal disjunctions) that we are able to create simulations of other 3-spaces more easily using loudspeakers than headphones. In both cases, however, there is a pressing need for intelligently designed proximity controllers which extend the current codifications for reverberation to enable the creation of more complex ambiences as described. The simple reverberator offering small or large concert halls and the like needs to retire and make way for programmable virtual space machines which react to sounds in much more complex and yet coherent ways.

These tools will become indispensable for the exploration of the immersive worlds of cyberspace in which, like in jungles, the point-of-hearing is omnidirectional, visual depth is very limited, and what surrounds us brings us in contact with the space. The

⁴ The separation of mind and the material world in such a way that the mind can only be acquainted with itself and the material world is known only through representation, effect of inference.

ear can lead the eye, the view being only revealed as we penetrate deeper into recalculated space.

10. CONCLUSION

The structuring rules of musical compositions are in some way codifications of sociopolitical ideologies, or they are at least ontologically related. They cannot be otherwise; not only through large-scale temporal structurings but in the unequal weighting given to one parameter over another in making well-formed utterances. In fact, the definition of some of the parameters themselves reflects this ontology.

Any experimental practice necessitates a renewal of experience with the materials used in order to develop new perceptions leading to new interrelationships and thus new methods of construction.⁵ Because the inner ear (the cochlea) is essentially a frequency analyser, perception of pitch is acute and this accounts for the primary role which it plays in music. However, we do not have sense organs for timbre and location. The perception of them is not only of a different order, but of a different kind. They are somehow 'constructed' or inferred from the *processing* of sensory information, yet we seem to experience these aspects of sound very directly, with little if any, conscious mental processing.

Timbre and 3-space have a number of commonalities. A functional approach to composition with both has some uses but suffers from its dependency, partly through notation, on static conceptualisations of what are in reality extremely multidimensional characteristics of sound. The need for these static conceptualisations for electroacoustic composition is weak especially when using morphological techniques. For musical grammars to be flexible enough for electroacoustic composition, they need to be able to break, or at least loosen these dependencies.

Just as Varese and others showed us how to hear a music in which pitch functions as a component of timbre, so I believe will we learn to use sonic proximity and radial location in a different way. Every sound we hear is a combination of the direct and reflected fluxions of a source in its environment. And if the sound is a recording of an environment, it includes the ambience of the recording as well as that of the listening environment. The nature of these sonic reflections – or filtrates – provides us with information from which the quite specific nature of the 3-space environment can be decoded. All sounds have

⁵ In music this happens most frequently with melody – the carrier of 'tone' – which observation explains Schoenberg's use of *Sprechstimme* and the 'spoken song' of rap music: new sonic structures outside the confines of the current theories of construction are discovered in the inflections of speech, thus encouraging a renewal of the language enabling new structural and formal possibilities.

the characteristics of the 3-space in which the source of the sound is emplaced, and the nature of these filtrates helps us to define the 3-space itself.

The need to hear afresh, outside the confines of 'music' which primarily uses sound as a means of carrying a message ('sonic rhetoric', or 'sound with attitude', if you like), is the reason that some of us prefer to be called sound artists rather than composers; we choose to make sound-art, and sound-sculptures rather than engage in overtly musical rhetoric. It is a continuum, however. By stepping outside the confines of what is currently thought of as musical discourse, composers can hope to renew the relationship between sound and music, and so to inject this relationship with more expressive power.

REFERENCES

- Begault, D. R. 1994. *3-D Sound for Virtual Reality and Multimedia*. Cambridge, MA: AP Professional.
- Boulez, P. 1971. *Boulez on Music Today*. London: Faber and Faber. Originally published as *Penser la musique aujourd'hui*, Paris 1963. Trans. Susan Bradshaw and Richard Rodney Bennett.
- Bregman, A. S. 1990. *Auditory Scene Analysis. The Perceptual Organisation of sound*. London: MIT Press.
- Chowning, J. M. 1971. The simulation of moving sound sources. *Journal of the Audio Engineering Society* **19**(1): 2–6.
- Dannenberg, R. B. 1993. Music representation issues, techniques and systems. *Computer Music Journal* **17**(3): 20–30.
- Gibson, J. J. 1966. *The Senses Considered as Perceptual Systems*. Boston: Houghton Mifflin.
- Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Goldstein, B. E. 1984. *Sensation and Perception*, 2nd edn. Belmont, California: Wadsworth Publishing Company.
- Grey, J. M. 1975. *An Exploration of Musical Timbre*. Center for Computer Research in Music and Acoustics, Department of Music Report No. STAN-M-2, Stanford University, February 1975.
- Handel, S. 1989. *Listening: An Introduction to the Perception of Auditory Events*. Cambridge, MA: MIT Press.
- Kolers, P. A. 1972. *Aspects of Motion Perception*. Oxford: Pergamon Press.
- Lombardo, T. J. 1987. *The Reciprocity of Perceiver and Environment. The Evolution of James J. Gibson's Ecological Psychology*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Stockhausen, K. 1997. Interview with Paul Bronowsky in *Access All Areas*, broadcast 4 May, ABC Television, Australia.
- Theile, G. 1986. On the standardisation of the frequency response of high-quality studio headphones. *Journal of the Audio Engineering Society* **34**: 953–69.
- Vennonen, K. 1994. A practical system for three-dimensional sound projection. In *Proc. of the Symp. on Computer Animation and Computer Music*. Canberra: Australian Centre for the Arts and Technology.
- Vennonen, K. 1995. *An Ambisonic Channel Card*. Graduate Diploma Thesis. Canberra: Australian Centre for the Arts and Technology.
- Wessel, D. L. 1979. Timbre space as a musical control structure. *Computer Music Journal* **3**(2): 45–52.
- Wishart, T. 1994. *Audible Design. A Plain and Easy Introduction to Practical Sound Composition*. York, UK: Orpheus the Pantomime Ltd.
- Wishart, T. 1996. *On Sonic Art*. Simon Emmerson (ed.). UK: Harwood Academic Publishers.
- Worrall, D. 1989. A music and image composition system for a portable multichannel performance space: a technical overview. *Chroma, Journal of the Australian Computer Music Association* **1**(3): 3–6.