International Conference on Music and Emotion, (Durham, UK) September 2009

The emotional contents of the 'space' in spatial music

Peter Lennox

Signal Processing and Applications Research Group (SPARG)

University of Derby, UK

p.lennox@derby.ac.uk

Tel: (0)1332 593155

Abstract

Human spatial perception is how we understand places. Beyond understanding *what* is *where* (William James' formulation of the psychological approach to perception); there are holistic qualities to places. We perceive places as busy, crowded, exciting, threatening or peaceful, calm, comfortable and so on. Designers of places spend a great deal of time and effort on these qualities; scientists rarely do. In the scientific world-view physical qualities and our emotive responses to them are neatly divided in the objective-subjective dichotomy.

In this context, music has traditionally constituted an item *in* a place.

Over the last two decades, development of "spatial music" has been within the prevailing engineering paradigm, informed by psychophysical data; here, space is an abstract, Euclidean 3-dimensional 'container' for events. The emotional consequence of spatial arrangements is not the main focus in this approach.

This paper argues that a paradigm shift is appropriate, from 'music-in-a-place' to 'music-*as*-a-place' requiring a fundamental philosophical realignment of 'meaning' away from subjective response to include consequences-in-the-environment. Hence the hegemony of the subjective-objective dichotomy is questioned. There are precedents for this, for example in the ecological approach to perception (Gibson). An ecological approach to music-as-environment intrinsically treats the emotional consequences of spatio-musical arrangement holistically. A simplified taxonomy of the attributes of artificial spatial sound in this context will be discussed.

The emotional contents of the 'space' in spatial music

Introduction and Context

Current technologies permit increasingly fine control of "three-dimensionality" in sound, in terms of phantom images' perceptible directionality from a specified listening position. It is possible to place images almost[1] anywhere on the surface of a sphere surrounding the listener. Image-movement controls are fairly rudimentary, descended from stereo panoramic potentiometers ("panpots") that vary inter channel relationships in amplitude (and sometimes phase-), thus manipulating listeners' interaural differences. Distance, range or proximity (whichever term is used to describe the separateness between perceiver and phantom source) is less well served, as are *changes* in proximity (approaching or departing). Synthesised sound fields can lack something related to "realism". It is also possible, using microphones such as Soundfield ™, Tetramic™ to capture a signal set that describes a full circular or even spherical sound field in a real environment. Whilst these microphones fall far short of the mathematically higher orders of resolution feasible through synthetic means, they can exhibit that elusive almost-tangible realism. Naturally, it is feasible to engineer hybrids where a "found sound field" can provide a container *place* within which synthesised or individually captured items can be sited, giving a holistic and natural-sounding environment yet with precisely locatable components.

Meanwhile, software for sound field control, scene-description protocols for audio and audio-visual spatial displays, along with storage and transmission technologies,

---

[1] Directly below the listener is still a problem

are evolving rapidly, driven by the growing ubiquity of domestic surround sound for film, TV and computer games.

Unsurprisingly, some composers are interested in notions of a paradigm shift using "space" as a direct musical parameter – not simply arranging musical sources around the listener, but utilising the full range of spatial attributes available to us in real environments. This raises many questions: is *space* musically meaningful at all? And if it is, how will space fit in with existing theories of music cognition? For that matter, what is "space"? Is it unitary or divisible? (If the former, it hardly seems possible to parameterise it, if the latter, what distinctions form the basis of parameterisation?)

For composers, the current situation is exciting and frustrating in equal measure. On the one hand, they might be eager to escape the constraints imposed by consumer technologies and formats; on the other hand, they are reliant on them for dissemination. With notable exceptions, music technologies for public concerts fall short of what is theoretically feasible (understandably; who would install a system for which no material exists, and who would compose for a similarly nonexistent system?). Particularly frustrating is that the detailed, exciting and no doubt magnificent spatio-musical imaginings are locked in the composer's mind, there is no intuitive and transparent conduit for externalising them and the media for their expression is currently "under construction"

The result is that spatial music must be laboriously and painstakingly constructed using tools that are not task specific; spontaneity, intuition and "feeling" are difficult to preserve in this way of working. This *can* result in a cerebral, analytical –even unemotional- approach to spatial composition.

This paper examines the philosophical and metaphorical constraints on 3 dimensional music development.

Artificial Space: intrinsic philosophies and metaphors

> Technology is at present covert philosophy; the point is to make it openly philosophical (Agre 1997 p.240)

Ultimately, *all* technologies feature embedded metaphors and are grounded on some axiomatic philosophical approach. When we press an onscreen 'button' we aren't really– we're issuing a 'command' to initiate a complex set of processing procedures… it's just 'dressed up' as a button to provide what Kahneman and Tversky called 'intuitive accessibility' (Kahneman 2002). When we turn a radio's volume *up* or *down*, the vertical spatial reference is metaphorical. It is interesting to note how many metaphors refer to spatial behaviours and the egocentric reference frame [Campbell 1993). For a detailed discussion of Metaphor Theory see (Lakoff and Johnson 1980, Lakoff 1993)

The underlying concepts of "space" might insidiously differ between one technological implementation and another. For instance, in domestic music technologies, the perceiver is conceptualised as a static, passive receiver, sitting quite still in a known position. This simplifies the signal control problem enormously, even though actually, very few perceivers are so fortunate or well behaved (in a car, no-one is *ever* seated in the 'correct' place). Computer gamers, on the other hand, tolerate quite poor-quality sound, but as long as the spatial audio aspects are consistent with their activities as a moving (virtually), behaving and exploring individual, they report an increased sense of 'presence' in the artificial environment.

Space: the engineering approach.

In engineering terms, space just objectively *is*, irrespective of my (or anyone else's) opinions. It is Euclidean, 3-dimensional and non-value-weighted. "Things" therein are measurable but not intrinsically *big* or *small*. Spatiality isn't quite the same kind of property as taste or colour (where the perceiver contributes a subjective element to the percept). To exist *is* to be spatial, and vice versa; a *thing* is spatially *extended* and *located*.

To engineer an artificial space, then, one has to control the smallest elements of which that space is made – the pixels, particles, waves, nodes or whatever- over the nominated display area that the percipient will experience. Finer control over smaller elements means "higher definition" with the concomitant assumption that the result will be more perceptually satisfying.

Spatial perception: an engineering problem

For the purposes of the present discussion, auditory spatial perception is conceived as itself a 'signal processing' problem, the perceiver apprehends signals (modulated energy flows) in just such a manner as to preserve the signal characteristics that physically correlate to the relevant[2] spatial attributes in the environment. The degree to which this is achieved is a matter of "spatial acuity", and of course, the degree to which there is a match between the resolution of a display system (the "high definition" of the previous paragraph) and the spatial acuity of the perceiver is the

[2] Relevant to cognitive processing and behavioural capacities and, supposedly, survival needs

degree to which a system has efficiently met requirements. In other words, higher-than-perceptible definition would be a waste of system resources.

This way of thinking could be simplistically blocked out as follows:

1. A sound field (characterisation of energy field at point of measurement) is presented to:
2. Sensory receptors, which sample the field, producing:
3. Transduction to neural analogue, representing key features prior to:
4. Cognitive processing to improve signal-to-noise ratio for salient features, then:
5. Higher cognition to understand meaning, attachment of emotional labels

This is a computational model of perception, where “meaning” is wholly intrinsic to the perceiver, is part of the *description* (of an external state of affairs) that comes into the equation at the end, after much computation – it is part of how we understand, through subjective weighting, aspects of our environment. This is classical Cartesian dualism wherein the interactions between physical universe and the mind are essentially mysterious and curiously indirect, mediated by layers of sensory and perceptual processing that progressively contaminate information with subjective self-interest.

In this conceptualisation, where value-weighting is assigned by the perceiver, the spatial distinctions *in* the environment are non-arbitrary, whereas those in perception are arbitrary. Hence, the fact that a perceiver can detect one feature but not another is attributed to inaccuracy in terms of acuity; perception is always less than perfect.

Unfortunately, artificial environments are similarly incomplete (by definition!) and so the spatial distinctions in this kind of environment must be arbitrarily parameterised to match perceivers' available distinctions. Engineers aim for "…as accurately as possible…" but it is really an impossible task, like chasing infinity.

Currently, artificial sound fields suffer the following constraints:

3. They are generally single room affairs and so not very 'explorable' (there are no very large artificial sound environments that perceivers could explore for hours or even days).

4. They are not nearly as richly detailed as real audible environments

5. Since we don't fully understand human spatial perception, we don't know how to compose for it, play with it, confound it.

6. There is little relevant *aesthetic* theory or practise of spatial sound

7. Empty, abstract, Euclidean space is dull and lifeless, rather like architectural drawings that have no people in them.

Many of these constraints come about simply because the technology has necessarily been driven (almost exclusively) by engineering concerns and problems, coupled with the spate of rapid evolution in digital technologies. Systems are complex to learn (and iron out the bugs) and quickly become redundant (often accompanied by the demise of the manufacturer and the cessation of technical support). Composers and musicians have had to become programmers, engineers, technicians before they can explore the aesthetic 'wild territory', subverting perceptual rules and driving technology forward with impossible demands. Listening to some of those composers who have thought about what spatial music might be like, I'm struck by the observation that the kinds of spatial behaviours they hear in their imagination and would like to incorporate are nothing like what is currently available. They speak of swarming, flying, bouncing coalescing, interacting, trajectories, call and response over great distances, huge spaces, waterfalls of sound, almost tangible yet abstract items moving very close then zooming far away and so on.

This is not to belittle the impressive results of the Cartesian engineering approach, not least because it's the only game in town, but given that the spatial characteristics available to composers must be arbitrary, a broader philosophical view will pay dividends.

Spatial perception: a problem for real people

An alternative approach casts spatial perception as a subset of *place perception*, whereby perception did not evolve to grasp abstract, empty space or objective reality or perfect Platonic forms; Euclidean space is an evolutionary latecomer to our cognition.

> "The doctrine that we could not perceive the world around us unless we already had the concept of space is nonsense. It is quite the other way round: We could not conceive of empty space unless we could see the ground under our feet and the sky above. Space is a myth, a ghost, a fiction for geometers." (Gibson, 1979 p. 3).

In Gibson's Ecological Approach, perception, cognition and action are indivisible. Cognition is *necessarily* embodied (see Lakoff 1993 for this view) so that perception and that-which-is-perceived (environment) are intertwined. See also Jarvilehto (1998) on the perceiver-environment *system*.

In this way of thinking, *real* places are intuitively accessible (as in Kahneman's 2002 description) precisely because the physical regularities and distinctions in them have shaped the species and the individual- our phylogenetic and ontogenetic development. "Perception" isn't inaccurate, subjective, biased, a poor-quality rendering of 'perfect' reality. Perception evolved (over millions of years) in a universe suffused with consequences: threats, opportunities. Meanings are not simply imaginary constructs, they exist in the environment. This question of what meaning *is*, is brought into focus in Victor Rosenthal's formulation of Microgenesis,

> It should be noted that form, meaning and value are not considered separate or independent entities. According to microgenetic theory, whatever acquires the phenomenological status of individuated form acquires, ipso facto, value and meaning. (Rosenthal, V., 2002)

The consequence of taking the view that meanings are not solely intrinsic to the perceiver is that we can accept that reasonably useful perception can take place when the signals available to sensation are in some way impoverished, or even when there is *so much* detail available that perception must somehow choose what to process *prior* to comprehension – the familiar "attention" problem. Perception is essentially non-linear, always must proceed on the basis of incomplete data (to arrive at good solutions in timely fashion), and the appropriate 'bits' of the incoming sensory signals must be rapidly sorted and assigned to a representational scenario that most resembles the important factors in the real world. This might seem an impossible task, since often, the 'bits' of the incoming sense-stream that represent different meanings in the world, are physically conflated. In hearing, for example, the spectral consequences of source *content*, and those due to source *location* are physically conflated in the signals at the eardrum – the 'what-and-where' problem. Nevertheless, we know that it is possible since we *do* achieve it.

Let us assume that all sensation carries information; the sensory field is always heterogeneous, patterned. Every single signal that impinges on a receptor is meaningful – but some meanings are more important than others, and *some* are more *urgent*. The supposition, then, is that there are natural hierarchies of *causal significance*, which are target items for mechanisms that analyse for *perceptual significance*. (Lennox and Myatt 2006) It is important to recognise that it cannot be necessary to understand the *entirety* of an environment to understand *any* of it. Hence, one could understand that *something* is moving without understanding *what* it is, or even *where* it is – these understandings can follow the raw recognition of something-moving.

Perception, then, is non-linear – and a good thing too. It has evolved to apprehend items of causal significance (meanings) in the world, by detecting and even *exaggerating* the physical aspects of incoming signals that correlate with those meanings. Perception is *primarily* identification and prediction; the conversion of sense data to meaningful-information-about-the-environment is crucial, but only a small part of that.

A prospective categorisation of items for hierarchical organisation based on the distinctions in real environments is as follows:

Entities, non-entities, events and relationships,

1. Entities are either *corporeal* or *ethereal*.

Ethereal*: 'an idea', 'a sound,' 'an event' 'a collection' or 'association'.*

Corporeal*: 'things'* or '*features'*.

Things*: organisms or objects*

Organisms: *Prey, predator, ally, competitor or crowd.*

Objects: *Tools, weapons, food*

Features *(*'entities of *potential* facility'*)*:

Obstacle, trap (self or other), shelter (hide behind/under), way (gap, doorway path, escape, gain access), vantage point (safe, remote viewing)

2. Non-entities

Formless Substance *e.g.: air*

(Perceptual) Background*: collection of uninteresting/non-urgent items (e.g.- grains of sand)*

Place: container *in which* action takes place

3. Events

Bounded: *sequence of changes involving cause, process and effect.*

Ongoing processes*: causal sequence without discernable boundary.*

4. Relationships:

Near / far: *Within reach/ can reach 'me', or not. Affects 'perceptual significance' of items.*

Nearer /further than *in comparison with other salient item(s)*

Moving: *Signifies entity, possibly animate.*

***Change* of movement** *–acceleration/ direction; good signifier of organism-hood. Predictions require frequent updates.*

Coming / going*: Threat or reward, imminent or receding*

Passing*: Salient change of significance, from '**coming'** to '**going'**.*

Facing-ness: *Characteristic of entities that facilitates prediction; can be a good signifier of organisms' intentions.*

(Lennox 2006)

It is obvious that these categories must interact; for example a large *thing* exhibiting organism-like behaviour, moving fast, towards me, potentially requires immediate action and is correspondingly likely to command attention in a way that causes me to ignore features of lesser urgency. Equally, in the absence of such urgent call-to-action, I can pay more detailed attention to very fine distinctions; for example, an interesting picture, a flower, a tiny almost-imperceptible sound.

Composing spatial music as an artificial environment

An important observation drawn from teaching in this subject area is that most people, most of the time, do not actually listen *to* sound environments. They hear *things* in places, near or far, behind or in front of something, coming, just moving or even stationary. Many of the items that figure in their perception are not represented in 'consciousness'; people don't *hear* them, they are *just there*. The obstacles such as buildings and trees that provide occlusions and reflective surfaces with which the sounds of everyday sources interact form a kind of unreported perceptual background. This is surely not a case of a perceptual deficiency; the deliberate ignoring of specific 'classes' of information in order to free up attention for others, is a complex accomplishment; I think of this as "selective inattention" (Lennox 2006). The causally less significant items and features in the world are suppressed until it seems that they are unheard – this is an example of the hierarchy of *perceptual significance* (Lennox Myatt and Vaughan 1999) matching the prevailing causal significance. In "normal" perception, we each do this apparently effortlessly, all day. It is after all, the faculty that has been shaped by many millions of years' exposure to the causal regularities of the world.

Nevertheless, when it comes to *examining* the constituents of real soundscapes, Kahneman's intuitive accessibility is problematic. I have had a group of postgraduate students insist that there are no echoes in a forest – that all acoustic energies are scattered and/or absorbed. This was because they took their knowledge from a textbook rather than a forest; a short trip to the countryside resolved the argument.

I sent a group of undergraduate students out on sound walks to listen to soundscapes. In spite of me giving detailed instructions on how to actually listen to the world rather than be deafened by one's own noise, most were reluctant to turn off their phone, it wasn't obvious to them that they shouldn't wear noisy clothing, or have their hood up. Some tried to walk around the route in a group, chatting, and one came back with an mp3 player on! Naturally, all were unimpressed by the soundscapes they had just experienced (or failed to experience). I sent the same students out in small groups with a Soundfield ™ ambisonic 3-D microphone and 4-track digital recorder, to record at various points along the route. We auditioned the results in our spatial sound lab which is suitably equipped with ambisonic decoders and 24 speakers arranged on the surface of a nominal sphere. They were astonished at how much the recording had picked up that they hadn't heard, how loud the background sound was and conversely, how indistinct some items were that they thought they had heard perfectly well. The students subsequently revisited the original soundwalk task and *almost* all returned surprised reports of being able to hear buildings, parked cars, passing openings and so on. The main impediments were said to be traffic noise and *fans* – extractor fans, heating fans, and air conditioning vents. It seems, as some have said (see, for instance: Schafer, 1969, and Truax 1996) our perceptual capacity to ignore the background comes at a price, we are generating the sonic equivalent of litter.

The lesson from the above experiences is that sometimes, abstraction is required, so that by rendering real soundscapes into the artificial domain, we can usefully gain information about the real world. Effectively, we are re-tuning and *elaborating* perception. This may the fundament reason for engaging in any abstract thought.

In supposing that an artificial environment must essentially be benign, the actual full range of causal significance (from inconsequential to immediate and severe threat) cannot be available; health and safety legislation and ethical considerations simply prohibit that. However, in the principle that perceptual significance can ‘scale’ to the immediate environment, the ‘dynamic range’ of perceptual experience is not necessarily curtailed; very fine distinctions can be entertained. Indeed, this must be the principle at work in our enjoyment of fictional environments such as music, film, reading and games.

Given that: *things*, *features* and *place* characteristics have causal attributes that cause perception to define them as such, it follows that their treatment in artificial sound fields must preserve those characteristics so that the environment exhibits an *ecology* consisting of items of more and less importance, changes in relationships, and causal progressions that accord with some “internal causal connectedness” that can facilitate perceptual anticipation (and occasional surprise!). That is to say, everything cannot be ultimately random at every scale, or the result will be perceptually uninteresting. Everything cannot be of maximal importance, nor can items be fundamentally unrelated.

This is not to say that spatial music must be slavishly spatially-realistic; we know from experience of cartoons and animated desk lamps (Pixar 1986) that literal realism is not necessary. Likewise, we know that spatial music will not all be the same, will not even feature the same entities, items or rules. However, inasmuch as there are rules, it would be useful if these were technologically implemented in transparent ways; this means that the control surfaces should metaphorically resemble the operation to be carried out.

The modular nature of modern digital audio technologies is useful, not least because currently, we tend to conceptualise perception itself as modular. That is, specialised processing subsystems 'concentrate' on certain types of task, feature or process. Without wishing to review the whole of the literature on brain region specialisation, Fodor's view on low-level modularity (Fodor 1985), 'what and where' processing systems (Atkinson 1993, pp 325-339), (Ungerleider and Mishkin 1982), the view espoused here is that *perceptual context modules* (Lennox 2006) are specialised to engage with understanding, tracking changes and generating predictions of specific causal features. So, for instance, a system that is "interested" in the intentions and behaviour of organisms would command attention (i.e. maximal processing resources] if an organism is nearby. A system that monitors "things coming toward me" (in vision this is referred to as *visual looming* (Franconeri and Simons, 2003), in audition, *auditory looming* (Rosenblum, Wuestefeld, and Saldaña, 1993) will cooperate with the intention monitoring system in sequestering attention-resources. Meanwhile, the system that subtracts background helps to improve matters by removing "noise" (i.e. unwanted signal) from the forum of attention; in audition, dynamic *precedence effect* mechanisms operate, so that the directional information linked to sonic reflections is suppressed in favour of a directional conclusion dominated by the first-received example of a sound: the 'direct' sound from the source itself (for a comprehensive description, see: Litovsky *et al*1999)

Hence, specialised spatial sound controls could be constructed so as to appeal specifically to evolved perceptual mechanisms for organism, object, place, movement, location in place, approaching, departing etc.

Plausibility rests with the constraints one places on the dedicated control. For instance, a human voice cannot come from several places at once, unless it is to be perceived as coming from a speaker system. Introducing small time delays around the multi-channel speaker array makes it sound as though one is sharing a reverberant environment with the person because these sound like acoustic reflections (about 20ms would be equivalent to a reflected sound wave having travelled 6.6 metres further than the direct path). It makes more sense to have a background "place management module" that gives the overall sound of a place of specified size, shape, reflectivity, clutter. Individual items can be positioned by a location module, or moved by a movement module. Things should move coherently and plausibly (few things can hop from one location to another without being detected en route, and a given thing can only move so fast and has finite acceleration). Many surround sound technologies currently feature circular panners; this is fine for choosing a particular location for a static source, but is ludicrous as a dynamic control. *Things* do not move in circles unless they are on the end of a piece of string or otherwise constrained, they tend to move in a straight line. The chief ubiquitous constraint that has shaped the environments we inhabit is *gravity*, so things bounce, bump, scrape and roll. Most things have mass and therefore *weight*. A thing that can maintain distance from the ground must be able to do so for a reason, and this reason should be perceptually discernable. Things do not start and stop moving for no reason – they are either moved by an agent or are agents themselves – and the movement in the two cases is likely to be different (irregular changes in direction and acceleration usually signify organic agency). Things do not inexplicably change size, come into or go out of being without explanation, things exhibit *continuity of identity*.

Of course, these don't sound like musical rules, they are the regularities of environments. Like the physics engines used in computer games, designing a 'perceptual engine' that utilises the intuitive and inbuilt distinctions we deploy in acts of spatial perception can simplify operations.

Developing a complete and prescriptive taxonomy of operations might not be useful at this stage, since the contention here is that the aesthetic 'rules' cannot be decided before the fact, and certainly cannot be decided without the creative input of composers and musicians. The selective confounding of perceptual expectations by varying parameters (so that a thing can dissolve into a hundred audible fragments, then coalesce into a different thing, for instance) is what spatial music will entail.

What is important is that music that the listener can be inside; this music-as-an-environment doesn't just immerse them in a thick blanket of sound, nor should it rely simply on arranging sources around a listener. This artificial environment should be able to engage the occupants' attention, inattention, anticipation and *feelings*. We must accept that the feelings we have about places are not merely due to subjective responses brought about by operant conditioning; places do have intrinsic character and emotional engagement is the *proper* response to place.

Conclusions

For best results, spatiality cannot simply be 'added on' to music (though doubtless that is how we get started). What is needed is to make explicit a coherent philosophy of the psychology of perception in an artificial spatial auditory environment. Sounds cannot exist in splendid isolation; they form part of an ecosystem of interdependent relationships. The tools for managing such an ecology barely exist, and their development cannot rest with engineers alone. Input is required from philosophers, psychologists and especially, artists.

In a small way, at the University of Derby we have initiated an attempt to bridge the gap between disciplines by creating the Derby Laptop Orchestra (DLO), which is a true 3-d sound environment for aesthetic, perceptual and technological experimentation. Whilst we work with, and contribute to the development of, some of the most advanced spatial sound techniques currently in existence, we are aware that we are at the very beginning of a huge and explosive growth, a veritable revolution that will soon make today's technologies seem like toys. The really impressive developments, though, will be in the aesthetics of this new medium.

We work with composers, psychologists, musicians, technologists, sound recordists, DJs, a professional audio company, geographers, architects, even a *historian* to explore matters of meaning in spatial sound. Naturally, whilst resources are necessarily limited, we welcome opportunities for further exploration.

References and Bibliography

Agre, P., *Computation and Human Experience*, Cambridge University Press 1997

Atkinson, J., (1993) "A neurobiological approach to the development of 'where' and 'what' systems for spatial representation in human infants" in (Eds.) Eilan, N., McCarthy, R., and Brewer, B. *Spatial Representation* Oxford. Oxford University Press.

Campbell, J. (1993) "The role of physical objects in spatial thinking" in: *Spatial Representation* Eds., Eilan, N., McCarthy, R., and Brewer, B. Oxford. Oxford University Press

Fodor, J. 1985. Precis of the Modularity of Mind. *Behavioral and Brain Sciences*, vol 8, pp. 1-5.

Gibson, J.J. (1979) *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston

Gibson E. J. (1988); "Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge" *Annual Review of Psychology* 39: 1-41

Jarvilehto T (1998). "The theory of the organism-environment system: I. Description of the theory". First pub. 1998 in *Integrative Physiological and Behavioral Science, 33, 321-334.* Retrieved August 2004 from: http://wwwedu.oulu.fi/homepage/tjarvile/orgenv1.pdf

Kahneman, D. (2002) "Maps of Bounded Rationality: a Perspective on Intuitive Judgement and Choice". *Nobel Prize Lecture* December 8, 2002

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago, IL: University of Chicago Press.

Lakoff , G. (1993) The Contemporary Theory of Metaphor in Ortony, Andrew (ed.) *Metaphor and Thought* (2nd edition), Cambridge University Press.

Lennox, P. and Myatt, T. (2007) Concepts of perceptual significance for composition and reproduction of explorable surround sound fields. In: Scavone, G.P., (ed.) *Proceedings of the 13th International Conference on Auditory Display*. ICAD2007, June 26-29, 2007, Montreal, Canada. , pp. 208-212.

Lennox, P.P., Myatt, A. and Vaughan, J., (1999) "From Surround Sound to True 3D". In: *Proceedings of Audio Engineering Society 16th International Conference*, AES, Rovaniemi, Finland,

Litovsky, R.Y., Colburn, H.S., Yost, W.A. and Guzman, S.J., (1999) "The precedence effect" in: *Journal of The. Acoustic Society of America* 106 (4), Pt. 1, October 1999 pp. 1633-1654.

Pixar Animated Film 1986 http://www.pixar.com/shorts/ljr/index.html

Rosenblum, L.D., Wuestefeld, A.P., and Saldaña, H.M (1993). "Auditory Looming Perception: Influences on Anticipatory Judgments". *Perception*. 22, 1467-1482

Rosenthal, V. (2002) "Microgenesis, immediate experience and visual processes in reading" in Carsetti, Arturo, Eds. *Seeing and thinking*. Kluver

Schafer, R. M. (1969) *The New Soundscape*. Vienna: Universal Edition

Truax, B., (1999). *Handbook for acoustic ecology* Second Edition (1999). ed. Truax, B., Cambridge Street Publishing. Retrieved Aug 2004, from: http://www2.sfu.ca/sonic-studio/handbook/Sound_Propagation.html

Ungerleider, L., and Mishkin, M. (1982) "Two cortical visual systems" In D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (Eds.), *Analysis of Visual Behaviour*, pp. 549--586. Cambridge, Mass. The MIT Press