What in the World Do We Hear?
An Ecological Approach to Auditory Event Perception

William W. Gaver
Rank Xerox Cambridge EuroPARC
61 Regent Street, Cambridge, CB2 1AB, U.K.

ABSTRACT

Everyday listening is the experience of hearing events in the world rather than sounds per se. In this paper, I take an ecological approach to everyday listening in order to overcome constraints on its study implied by more traditional approaches. In particular, I am concerned with developing a new framework for describing sound in terms of audible source attributes. An examination of the continuum of structured energy from event to audition suggests that sound conveys information about events at locations in an environment. Qualitative descriptions of the physics of sound-producing events, complemented by protocol studies, suggest a tripartite division of sound-producing events into those involving vibrating solids, gasses, or liquids. Within each of these categories, basic level events are defined by the simple interactions that can cause these materials to sound, while more complex events can be described in terms of temporal patterning, compound, or hybrid sources. The results of these investigations are used to create a map of sound-producing events and their attributes useful in guiding further exploration.

INTRODUCTION

Imagine that you are walking along a road at night when you hear a sound. On the one hand, you might pay attention to its pitch and loudness and the ways they change with time. You might attend to the sound's timbre, whether it is rough or smooth, bright or dull. You might even notice that it masks other sounds, rendering them inaudible. These are all examples of *musical listening*, in which the perceptual dimensions and attributes of concern have to do with the sound itself, and are those used in the creation of music. These are the sorts of perceptual phenomena of concern to most traditional psychologists interested in sound and hearing.

On the other hand, as you stand there in the road, it is likely that you will not listen to the sound itself at all. Instead, you are likely to notice that the sound is made by an automobile with a large and powerful engine. Your attention is likely to be drawn to the fact that it is approaching quickly from behind. And you might even attend to the environment, hearing that the road you are on is actually a narrow alley, with echoing walls on each side.

This is an example of *everyday listening*, the experience of listening to events rather than sounds. Most of our experience of hearing the day-to-day world is one of everyday listening: we are concerned with listening to the things going on around us, with hearing which are important to avoid and which might offer possibilities for action. The perceptual dimensions and attributes of concern correspond to those of the sound-producing event and its environment, not to those of the sound itself. This sort of experience seems qualitatively different from musical listening, and is not well understood by traditional approaches to audition.

The distinction between everyday and musical listening is between experiences, not sounds (nor even psychological approaches). It is possible to listen to any sound either in terms of its attributes or in terms of those of the event that caused it. For instance, while listening to a string quartet we might be concerned with the patterns of sensation the sounds evoke (musical listening), or we might listen to the characteristics and identities of the instruments themselves (everyday listening). Conversely, while walking down a city street we are likely to listen to the sources of sounds – the size of an approaching car, how close it is and how quickly it is approaching – but occasionally we might listen to the world as we do music – to the humming pitch of a ventilator punctuated by a syncopated birdcall, to the

---

Gaver, W. W. (1993). What in the world do we hear? An ecological approach to auditory source perception. *Ecological Psychology, 5* (1): 1-29. Comments to w.gaver@rca.ac.uk.

interplay and harmony of the sounds around us. This may seem an unusual experience, but hearing the everyday world as music is one way to understand what John Cage, for instance, is attempting in his compositions. In presenting traffic sounds, for example, in a concert setting, he is trying to evoke an experience of musical listening to non-musical sounds (see Cage, 1976).

The Psychology of Everyday Listening

The experience of everyday listening may serve as the foundation for a new explanatory framework for understanding sound and listening. Such a framework would allow us to understand listening and manipulate sounds along dimensions of sources rather than sounds. So for example, we might imagine a psychoacoustics concerned with measuring people's ability to hear the forces involved in events, or a synthesiser that allowed us to specify whether a sound source was wood or metal.

Such an account would complement the one offered by traditional psychoacoustics. Clearly the phenomena studied in traditional psychoacoustics are valid whether a person is engaged in everyday or musical listening: A loud sound will mask a soft one, whether I am concerned with the soft one having a certain pitch or with the size of its source. But a new framework may alter some of the questions about sound and hearing we consider important. It certainly should allow us to refer more directly to attributes of everyday listening than does our present understanding of psychoacoustics.

For despite centuries of study of sound and hearing, we do not really know how it is, for instance, that a listener can hear that a sound is made by a large and powerful automobile approaching in a narrow alley, or can tell that somebody is ascending or descending a staircase from the sounds made. Laboratory studies concerning pitch perception, loudness, and so forth are as distant from these sorts of question as studies of peoples' perception of letter shape and spacing would be from understanding reading comprehension. Although research on audition has included topics relevant for everyday listening – most notably localisation and timbre perception – we still know little about how to characterise the fundamental attributes of events that we hear, or about the acoustic cues that inform us about these events and their attributes.

Why Is So Little Known About Everyday Listening?

There are two reasons for our ignorance, one historical and the other theoretical. Historically, studies of acoustics and psychoacoustics have been guided largely by a concern with understanding music and the sounds produced by musical instruments. From the ancient Greeks' discovery that doubling the length of a vibrating string lowers its pitch by an octave, through to Helmholtz' (1885/1954) studies of the harmonic structure of musical sounds, and even to present day studies of computer music (e.g., Pierce, 1983), the major thrust of disciplines concerned with nonspeech audio has been to use musical sounds and to understand musical listening.

But an account of hearing based on the sounds and perceptions of musical instruments often seems biased and difficult to generalise. Musical sounds are not representative of the range of sounds we normally hear. Most musical sounds are harmonic; most everyday sound inharmonic or noisy. Musical sounds tend to have a smooth, relatively simple temporal evolution; everyday sounds tend to be much more complex. Musical sounds seem to reveal little about their sources; while everyday sounds often provide a great deal of information about theirs. Finally, musical instruments afford changes of the sounds they make along relatively uninformative dimensions such as pitch or loudness, while everyday events involve many more kinds of changes – changes that are often musically useless but pragmatically important. Our current knowledge about sound and hearing has been deeply influenced by the study of a rather idiosyncratic subset of sounds and sources.

Studies of audition have been further constrained by sensation-based theories of perception and the supposed primitives of sound they suggest. Physical descriptions of sound are dominated by those suggested by the Fourier transform: frequency, amplitude, phase, and duration. Traditional explanations of psychophysics take these "primitive" physical dimensions as their elemental stimuli and use them to motivate the identification of corresponding "elemental" sensations. From this perspective, more complex perceptions must depend on the integration of elemental sensations – but often, sensations seem inadequate to specify complex events. Thus traditional approaches argue that there is often a paucity of information in available stimuli, and that veridical perception must depend on representations of the world based largely on memory, "unconscious inference" or problem-solving. The upshot of this approach is a strategy preoccupied with elemental sensations, while questions concerning auditory event perception – if recognised at all – are left to higher-level cognitive accounts. But because students of cognition are often more interested in general questions of mechanisms and representations than in the content of cognition, focused examinations of everyday listening often fall between the cracks separating psychoacoustics and cognition.

The Ecological Approach To Perception

In order to understand everyday listening, it is necessary to broaden the range of physical parameters and perceptual experiences to be considered. For instance, we might add new perceptual dimensions such as size or force to those considered by psychoacoustics, and study them in terms of their acoustic correlates. Such an endeavour implies that traditional psychoacoustics needs to be stretched in two ways: First, the perceptual dimensions we need to study concern those of sources as well as sounds, and second, we must be willing to treat apparently complex acoustic variables as elemental.

This approach is inspired and guided by the ecological approach to vision developed by Gibson (1979/1986, 1966). The ecological perspective counters many of the assumptions of traditional accounts of perception. According to the ecological approach, perception is usually of complex events and entities in the everyday world. Moreover, it is *direct*, unmediated by inference or memory. Elemental stimuli for perception do not necessarily correspond to primitive physical dimensions but may instead be specified by complex *invariants* of supposedly primitive features. Thus complex perceptions rely on seemingly complex stimuli (or “perceptual information”), not on the integration of sensations. From this point of view, there is rich and varied information in the world, both because our descriptions are no longer limited to primitive physical dimensions and because *exploration* of the world over time – as opposed to passive exposure to discrete stimuli – becomes an important component of perception. Thus, according to the ecological account, the study of perception should be aimed at uncovering ecologically relevant dimensions of perception and the invariant perceptual information for them.

Developing An Ecological Account Of Listening

There has been little development of an ecological account of audition as yet. Gibson (1966) did make some preliminary observations about listening, but did not develop them to a great degree. More recently, there have been several studies of listening based on the ecological perspective (e.g., Heine & Guski, 1991a, 1991b; Shaw et al., 1991; Gaver, 1988; Warren & Verbrugge, 1984; Vanderveer, 1979); these will be discussed throughout this paper and its companion (Gaver, in preparation). Such studies have proven informative on their own and lend support to the idea that an ecological approach to audition might be fruitful, but a comprehensive account of everyday listening has yet to emerge. There is little in the way of an organising theoretical framework that allows the relations among various research results to be understood, or that guides the formulation of new questions.

What might such a framework look like? It must answer two simple but fundamental questions. First, in expanding upon traditional accounts of elemental sensations, we must develop an account of ecologically relevant perceptual entities: the dimensions and features of events that we actually obtain through listening. Thus our first question is *What do we hear?* Similarly, in expanding traditional accounts of the primitive physical features of sound, we must seek to develop an *ecological acoustics*, one which describes the acoustic properties of sounds that convey information about the things we hear. Thus our second question is *How do we hear it?*

In the rest of this paper, I explore the first of these questions with the aim of developing an descriptive framework of the perceptual attributes and dimensions that characterise the auditory perception of events. Such a framework is useful both clarifying the relations among various projects concerned with everyday listening and ecological acoustics and in suggesting new avenues for exploration. In this discussion, I start by briefly considering the range of information about the world that we hear, and then focus more specifically on audible attributes of source events. In a companion paper (Gaver, in preparation), I describe a number of synthesis algorithms that allow sounds to be specified in terms of the events that cause them. For instance, instead of specifying a particular waveform modified by some amplitude envelope, one can request the sound of an 8-inch bar of metal struck by a soft mallet. These algorithms serve as instantiated theories both about what we hear and particularly about the acoustic information characterising various of the audible event attributes described here. Between the two papers, then, I hope to point towards a foundation for future work on everyday listening.

FROM EVENT TO EXPERIENCE

The example of hearing an approaching automobile is useful in understanding what we hear, and thus the scope of an ecological approach to audition. Figure 1 shows this event in the form of a continuum of energy between the source event (the automobile) and the experience. There are several landmarks along the way, each with its own characteristics as a medium for structured patterns of energy. Describing the route from event to experience in detail would involve volumes on mechanical physics, fluid dynamics, acoustics, anatomy, and psychoacoustics, and even then there would be much we do not know. Nonetheless, a brief examination of the continuum is useful for generating intuitions about what sorts of information might be available for listening (a detailed, though more traditional, account may be found in Handel, 1989; see Heine & Guski, 1991 for a review of this book from an ecological perspective). In particular, such an examination suggests that a given sound provides information about an *interaction* of *materials* at a *location* in an *environment*.

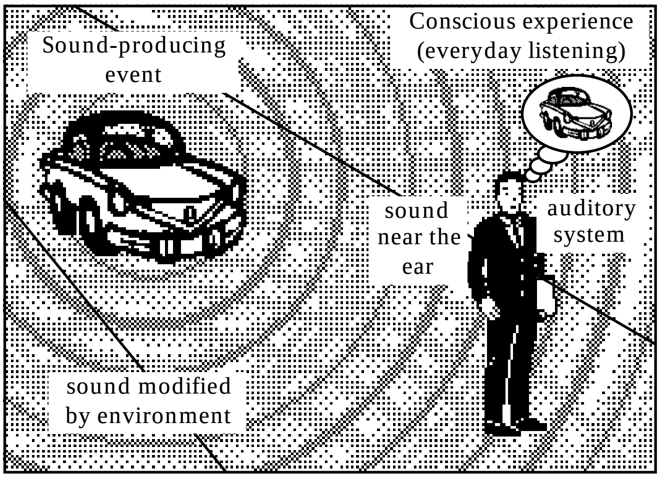

Figure 1: The continuum from world to experience. A source event causes sound waves: some radiate directly to an observation point, others are modified by the environment before reflecting to the listener. Invariant patterning of the acoustic array near the ear is picked up by the auditory system, providing information for the experience of everyday listening.

The first landmark is that of a *sound-producing event*, or *source* of sound. In the case of the automobile, some proportion of the energy produced by burning petrol causes vibrations in the material of the car itself (instead of contributing to its gross movement). Things tap, scrape, slosh, rub, roll, and flutter. Because each of these events is determined by the physical attributes of its source, the entire pattern of the car's vibrations is meaningfully structured by its components. These mechanical vibrations, in turn, produce waves of alternating high and low pressure in the air surrounding the car. The pattern of these pressure waves follows the movement of the car's surfaces, so that faster movement forms smaller regions of pressure differences, more complex movement forms more complex density patterns, and more extreme movement forms higher-amplitude variations of pressure (within limits determined by the frequency-dependent coupling of the surface's vibrations to the medium; see Handel, 1989). These spreading pressure waves, then, may serve as information about the vibrations that caused them, and thus for the event itself. When they change with sufficient amplitude and within a range of frequencies to which the auditory system is sensitive, the result is a sound wave from which a listener might obtain such information.

Each source of sound involves an *interaction* of *materials*. For instance, when two gears rub against each other, the patterns of their vibration depend both on the force, duration, and changes over time of their interaction as well as the size, shape, material, and texture of the gears themselves. The pressure waves produced by the resulting vibration are determined by these attributes, and thus may serve as information about them. Unlike radiant light, which is relatively unstructured and thus uninformative (in contrast to reverberating light), radiant sound is richly structured by – and thus may provide information about – its source. Describing what we hear of these sources is a primary goal of this paper.

Sound is also structured by and informative about the *environment* in which the event occurs. This is the second landmark in the continuum from event to experience. As Figure 1 indicates, much of the sound that reaches us from a source has reflected off various other objects in the environment, which colour the spectrum of reflected sound just as light is coloured by the surfaces it strikes. Because radiant sound waves always reach the ear before reflections or not at all (as when only echoes reach an observation point), information about sources may be distinguished from that about the environment (Hafter & Buell, 1985). In addition, the medium itself shapes the sounds it conveys: sound waves lose energy, especially high-frequency energy, as they travel through the air, and thus provide information about the distance of their sources. As sources move with respect to a potential observation point, their frequencies shift, producing the Doppler effect. Finally, changes in loudness caused by changes in distance from a source may provide information about time-to-contact in an analogous fashion to changes in visual texture (Shaw et al., 1991).

Sound converges on a potential listening point from every direction; the differences in intensity and spectrum in each direction provide an informative structure. The result is an *auditory array*, analogous to the optical array. Although the information conveyed by the auditory and optical arrays is different, and different mechanisms have evolved to pick up this information, the pattern of similarities and differences in different directions is informative for listening just as it is for looking. Reverberant sound is informative about the size and layout of the surrounding environment. We know this information is used by bats and blind people (Jenkins, 1985), it is likely that sighted individuals are sensitive to it as well.

One implication of the auditory array is that sound is not only informative about its source and environment, but about its *location* as well. A great deal of research has explored people's ability to localise sound sources and the information they use to do so (e.g., Guski, 1990; Wenzel et al., 1991; Lehnert & Blauert, 1991; Blauert, 1983). Directional information can be picked up because we have two ears, and each ear has a complex, direction-dependent filter (the pinna). This allows sensitivity to the timing differences and spectral shifts that specify direction. The pattern of radiant and reflecting sound can further aid localisation. Moreover, the auditory system is mobile. We can turn our heads to pick up shifting inter-aural patterns and to orient towards a sound, further improving localisation. Finally, Hoeger (1992) has explored the possibility that aural texture – the patterns of intensity and spectra of *groups* of sources at different distances – may provide further information about the layout of multiple sources in the environment.

Sources, Locations, and Environments

In sum, then, sound provides information about an *interaction* of *materials* at a *location* in an *environment*. We can hear an approaching automobile, its size and speed. We can hear where it is, and how fast it is approaching. And we can hear the narrow, echoing walls of the alley it's driving along. These are the phenomena of concern to an ecological approach to perception.

Traditional psychologists do not have much to say about this sort of information because they only study part of the continuum from source to experience. Typical research focuses on the sound itself, analysing it in terms of properties such as amplitude and perceived loudness, or frequency and perceived pitch. Such research misses the higher-level structures that are informative about events. Taking an ecological approach, on the other hand, implies analyses of the mechanical physics of source events, the acoustics describing the propagation of sound through an environment, and the properties of the auditory system that allow us to pick up such information. The result of such analyses will be a characterisation of acoustic information about sources, environments and locations which can be empirically verified. This information will often take the form of complex, constrained patterns of frequency and amplitude which change over time: These patterns, not their supposedly primitive components, are likely to provide listeners with information about the world.

In what follows, I focus on sound-producing events with the aim of describing their audible dimensions and attributes. Clearly this is only part of the story that an ecological approach to audition must tell. But it is an intriguing part, and one that is perhaps the most neglected. We know very little indeed about what we hear of the world, about the attributes of events that are specified by sounds. The rest of this paper is concerned with such an analysis.

THE PHYSICS OF SOUND-PRODUCING EVENTS

Sounds are determined by and informative about attributes of their sources. This means that audible source parameters – perhaps size and material, for instance – can be added to the list of dimensions of sound to be studied by psychologists. But which attributes of a sound-producing event actually affect the sounds it makes? Which are likely to be audible? In this section, I approach these questions by discussing the physics of sound-producing events in a qualitative way. The purpose here is not to provide a detailed account of mechanical physics (see, e.g., Skudrzyk, 1968, or Cremer, 1984 for such accounts.). Instead, the aim is to provide an initial orientation towards the relevant attributes of sound-producing events such as:

- closing a door,
- scraping fingernails over a blackboard,
- water dropping into a pool,
- wind whistling through wires,
- an exploding balloon,

- a resonating tuning fork.

It will become clear that these events share a number of common features – the most general being that all are caused by the interaction of materials. But there are substantial differences in the physics of closing a door, for instance, water dropping into a pool, and an exploding balloon. Most fundamentally, these events fall into three categories: those in which sounds are produced by vibrating solids; those in which sounds are produced by changes in the surface of a body of liquid; and those in which sounds are directly introduced into the atmosphere by aerodynamic causes. Here I first describe vibrating solids in some detail, and then use that discussion to illuminate liquid and aerodynamic sounds.

Vibrating Solids

Vibrating solids are a common source of radiant sound waves. This class of event includes slamming doors and scraping fingernails, footsteps and breaking glass, as well as the automobile featured in the description above of the continuum between event and audition. In this section I discuss the physical properties that affect the vibration of a solid object in order to suggest potentially audible attributes (see also Skudrzyk, 1968; Rayleigh, 1877/1929).

Figure 2 shows a simple example of a vibrating solid. Objects vibrate when a force is exerted upon and then removed from a system that is otherwise at equilibrium. This input of energy deforms the object from its original configuration; the forces that resist this deformation result in the build up of potential energy in the new configuration (Figure 2A). When the deforming force is removed, the object starts to return to its original shape, due to various *restoring forces* acting with the potential energy stored by the deformation (Figure 2B). This results in the movement of the object towards its initial configuration. But when it reaches this initial position, the potential energy has been converted to kinetic energy, and the object moves through the resting configuration (Figure 2C). Just as a pendulum swings back and forth once it is set in motion, the repeated translation from potential energy to kinetic energy and back again causes the object to move repeatedly through its resting configuration: It vibrates.

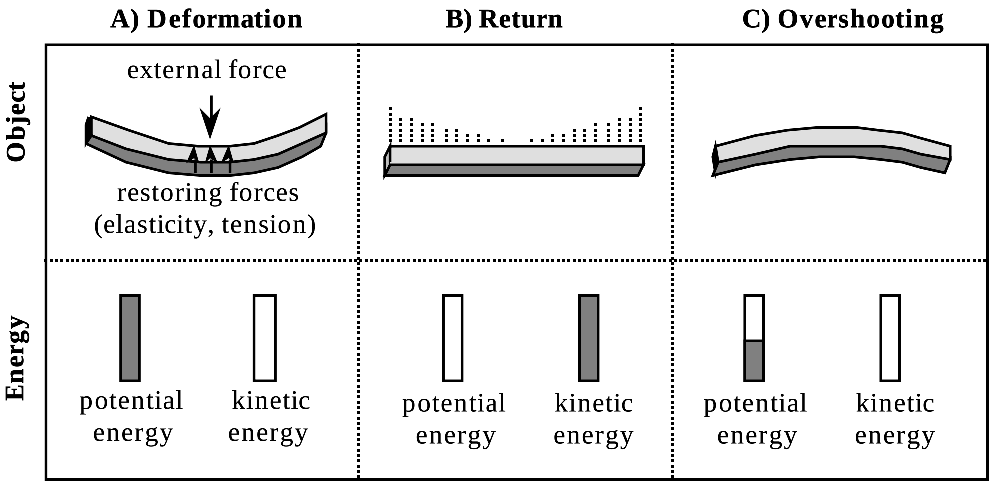

Figure 2. When an object is deformed by an external force, internal restoring forces cause a build up potential energy (A). When the external force is removed, the object's potential energy is transformed to kinetic energy, and it swings through its original position (B). The object continues to vibrate until the initial input of energy is lost to damping (C).

If no energy were lost in this translation, vibration would go on forever, and the world would be a very noisy and shaky place. But energy is lost for various reasons, collectively referred to as *damping* (see the energy measures in Figure 2C). The object moves until all the initial energy responsible for its deformation has been lost, and it returns to its old (or finds a new) equilibrium configuration.

The pattern of vibrations of a given solid is structured by a number of its physical attributes. Table 1 shows these properties grouped in terms of attributes of the *interaction* causing the vibration, those of the vibrating object's *material*, and those of the object's *configuration*.

Table 1

Source	**Effects on the Soundwave**
Interaction	
Type	Amplitude function, spectrum
Force	Amplitude, bandwidth
Material	
Restoring Force	Frequency
Density	Frequency
Damping	Amplitude functions; also frequency
Homogeneity	Complex effects on amplitude; also frequency
Configuration	
Shape	Frequency, spectral pattern
Size	Frequency, bandwidth
Resonating Cavities	Spectral pattern
Support	Amplitude functions, frequency, spectrum

The interaction *type* (e.g., whether the object is hit, scraped, rolled) affects both the time-varying amplitude and the spectrum of vibration. Impacts, for instance, involve a rapid and discrete deformation of an object, thus the amplitude of ensuing vibrations depends only on the damping of the system. Scraping, in contrast, involves a continuous input of energy into the system, which affects the amplitude of vibration over the duration of the event. The *force* of the interaction (i.e., how hard it is hit) determines the overall amplitude of vibration: when an object is struck or scraped with more force, it vibrates with greater amplitude (and makes a louder sound). Force also affects the bandwidth of vibration by deforming objects more sharply and thus increasing the high frequency components of subsequent vibration. In general, different sorts of interactions produce a wide range of characteristic amplitude and vibration patterns; it is probable that these patterns specify the events that cause them.

A number of material attributes also determine the patterns of vibration of a given system. The strength of the *restoring forces* (e.g., tension, elasticity) determines the potential energy implied by a given deformation, while the material's *density* determines its inertia; both together determine how quickly it will return from its deformed state and thus the frequency of its vibrations. *Damping*, the loss of energy from vibration, has various causes including internal heat transfer, plasticity, and external absorption of energy (including by the air as sound). The kind and amount of damping has clear effects on the overall time-course of vibration. For instance, the damping of wood tends to be much greater than that of metal, which is why wood "thunks" and metal "rings." Finally, the *internal structure* of the material has many complex effects on its vibration, particularly in the temporal domain. For example, the vibrations made by wood, which is both inhomogeneous and anisotropic, seem to damp out in a more irregular fashion than those of homogeneous and isotropic materials like metal (Gaver, 1988). In sum, material has a number of salient and characteristic effects on vibration which further specify sound-producing events.

The configuration of an object (e.g., its shape and size) also affects the sounds it produces. *Size* determines the lowest-frequency vibrations that an object can make: big objects tend to make lower sounds than small ones. The *shape* of an object determines the frequency and spectral patterns produced by objects. Bars, for instance, tend to make inharmonic sounds with widely-spaced partials, while the spectra produced by plates tend to be denser. These different spectral characteristics are further influenced by the existence and characteristics of any *resonating cavities*. Vibrations in the air of such cavities reinforce or damp corresponding vibrations in the surrounding material; instruments such as guitars or violins incorporate resonating cavities to amplify and colour the sounds they produce. Finally, the object's *support* may produce large effects in its vibrations, changing their spectral pattern as well as their frequency-dependent damping. A good example of the effects of support can be found by listening to windchimes: Most windchimes are hung near the top, which tends to damp out lower-frequency modal vibrations, leaving only the inharmonic and short-lived higher partials. Hanging windchimes by a point about a third of the way down their length (corresponding to the node of their fundamental modal vibration) allows them to make lower, longer-lasting sounds. In sum, then, the patterns of vibration of a solid object are determined by its configuration just as by its interaction and material; again we may suspect that these properties are specified by the sounds that they make.

The parameters discussed above may also be grouped according to their effects on the *frequency* or *temporal* domains (Table 2) – that is, whether they effect the frequencies with which the object vibrates or the changes of vibration over time. Of course, frequency is the reciprocal of time, so these domains are not physically different, but rather the result of different representations. But the two domains are separable psychologically – when things cause pressure variations that change quickly enough, they are perceived as a sound, otherwise as a change of sound – and thus the distinction is appropriate for an ecological acoustics.

Table 2

Frequency domain	**Temporal domain**
Restoring force	Interaction type
Density	Damping
Size	Internal Structure
Shape	Support
Support	

Attributes which make their effects in the frequency domain influence an object's initial return from deformation, and thus the frequencies of its subsequent vibrations. Attributes which influence the temporal domain produce effects that become apparent only after repeated cycles of vibration. In general, attributes of the object (e.g., the strength of restoring forces, density, size) tend to influence the sounds in the frequency domain, while attributes of the interaction (e.g., its type and force) tend to influence the temporal domain. While this correspondence is by no means perfect – for instance, the force of interaction can affect a sound's bandwidth in the frequency domain, and the damping of a material is a strong determinant of the sound's temporal behaviour – it is good enough to lend some support to Vanderveer's (1979) hypothesis that *interactions affect the temporal domain* of sounds, and *objects the frequency domain* .

Vibrating solids as described above include many common sources of sounds, such as footsteps, scraping blackboards, clattering silverware, closing doors, and so on. For instance, folding paper makes sounds for similar reasons as a struck bar. But when paper is folded, it doesn't return to its initial configuration but instead deforms along a line of stress, finding a new equilibrium configuration. The sound it makes is thus a result both of the sudden fold and the vibrating surfaces it separates.

This level of description can also serve as a foundation for describing other more complicated events such as crumpling paper (multiple deformations), breaking and bouncing (multiple impacts; see Warren and Verbrugge, 1984), and sawing wood (multiple scrapes). Though new source attributes may come into play in such events, those described here remain important. In general, it may be expected that these attributes are salient in determining the sounds produced by all events involving vibrating solids.

Aerodynamic Events

The properties of aerodynamic events are somewhat different than those describing vibrating solids. Where the sound waves created by solid objects are due to the interaction of a vibrating surface with the atmosphere, aerodynamic sounds are caused by the direct introduction and modification of atmospheric pressure differences from some source.

The simplest aerodynamic sound is exemplified by an exploding balloon (see Figure 3.A.) When a balloon bursts, a mass of high-pressure gas is released into the surrounding atmosphere. This sudden pressure variation propagates as a wave which may be heard if the pressure differences that reach the ear are large enough and if they change at an appropriate rate. In such events, sound is caused directly by sudden pressure variations in the air, not by the effects of a vibrating surface. Most of the information conveyed by explosions seem to be carried by the frequency bandwidth of the sound, and seems likely to concern the size or force of the explosion. High frequency components seem to indicate the suddenness of the pressure change near the source, while low frequency components depend on the amount of gas involved (and thus the duration of the initial pressure release). So one might hear large, sudden explosions, or smaller, less abrupt bursting noises.

A) Explosion

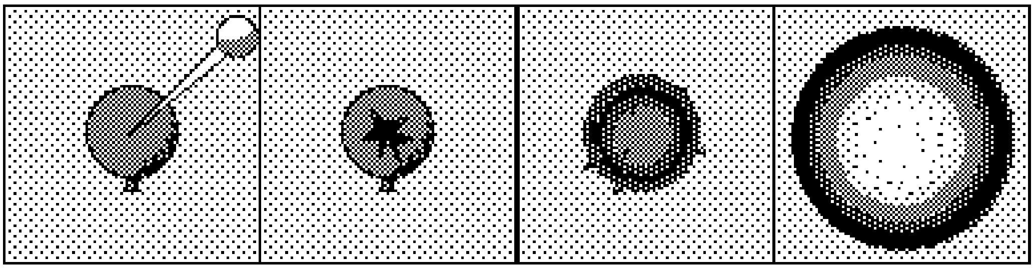

B) Flow Past a Cylinder

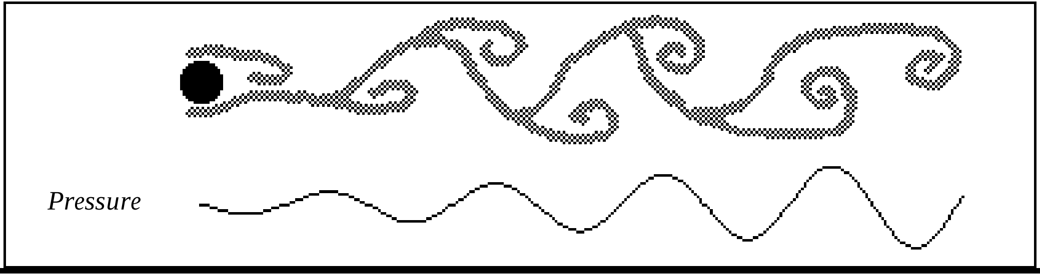

Figure 3. Aerodynamic sounds involve abrupt changes of air pressure, as when a balloon explodes (A), or wind rushes past a cylinder (B).

The sudden change in pressure caused by a bursting balloon or explosive is analogous to a discrete interaction (such as an impact) that causes an object to vibrate. Other aerodynamic sounds, such as the hissing of a leaky pipe or the rush of wind from a fan, are caused by more continuous events. The sounds made by these events are also caused by the direct introduction of pressure variations in the atmosphere, and are determined by various attributes of their sources. The sounds produced by leaky pipes, for instance, are determined by the pressure within the pipe and variations in this pressure caused by turbulence. The sounds made by wind rushing from a fan are affected by the speed and size of the fan, and thus the volume of air it moves. Finally, the density of the gases involved also affect the sounds, though listeners are likely to be relatively insensitive to this information. In general, we expect aerodynamic sounds to be made by air, which is by far the most common medium that we experience. That this is so can be seen by considering why the sound of someone talking after inhaling helium is so humorous. The heightened pitch of the voice is perfectly predictable due to the lower density of the gas, but quite unexpected on the basis of experience.

Another sort of aerodynamic event involves situations in which changes in pressure themselves impart energy to objects, causing them to vibrate. For example, when wind passes through a wire, eddies form on alternating sides, and the variations of pressure on each side causes the wire to vibrate (Figure 3.B; see also Van Dyke, 1982). The frequencies of vibration produced depend on the wind-speed, size, and tension of the wire; this is the principle used in creating Aeolian harps. In addition, sound itself may impart energy to objects, as when its minute pressure variations match the modal frequency of a tuning fork, causing it to ring through sympathetic vibration. Such sounds are not purely aerodynamic; instead, they involve an aerodynamic *cause* of a vibrating solid – though information may be available in such sounds about attributes of the gasses involved, just as it is about many of the interactions that cause sound.

Liquid Sounds

Sound-producing events involving liquids (e.g., dripping and splashing) are like those of vibrating solids in that they depend on an initial deformation that is countered by the material's restoring forces. But it seems the resulting vibration of the liquid does not usually affect the air in direct and audible ways. Instead, the resulting sounds are determined by the formation and change of resonant cavities in the surface of the liquid.

This can be seen most clearly in considering how an object dropping into liquid makes a sound, as shown in Figure 4. As the object hits the liquid, it pushes it aside, forming a cavity that resonates to a characteristic frequency, amplifying and modifying the pressure wave formed by the impact itself. The cavity grows as the object displaces more liquid, and thus the resonant frequency decreases. But the liquid's pressure causes it to close in on the cavity, decreasing the size of the resonant cavity, until the object is ultimately immersed and the sound stops. Meanwhile, as the object enters the body of liquid, a "crown" often forms from the displaced liquid around the cavity which disintegrates into smaller droplets each of which creates a small resonating cavity as it falls. Similarly, a "spout" may form in the middle of the closing cavity as the object is immersed, which again produces a sound as it falls (for beautiful photographs of these phenomena, see Edgerton & Killian, 1975). Thus a dripping sound is likely to be characterised by a short-lived pitched impulse followed quickly by several higher-pitched, even shorter-lived ones. The details of such sounds are likely to be influenced by many factors, particularly the mass, size and speed of the object and the viscosity of the liquid, all of which influence the evolution of the resonating cavity and the formation and separation of the crown and spout.

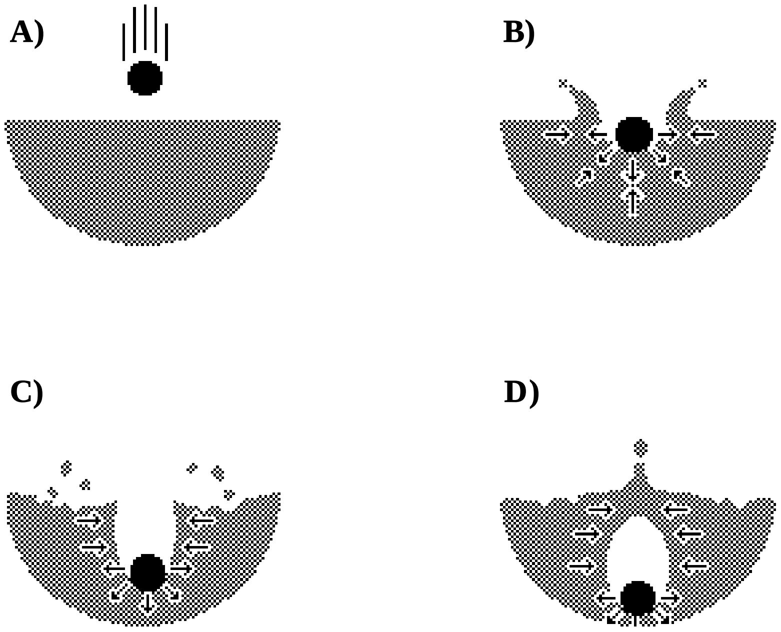

Figure 4. When an object falls into a liquid (A), it forms a resonant cavity with a characteristic frequency (B), which changes as more liquid is pushed aside (C). Finally, the liquid's pressure causes the cavity to close in on the cavity (D) until the object is completely immersed. Meanwhile, the displaced water forms a crown (B) which falls back into the water (C), and as the water rejoins it often forms a spout (D); both crown and spout contribute to the resulting sound.

More complex splashing sounds also seem to produce their sounds as changing cavities are formed which resonate, amplify and modify the sounds made by impacts of liquid on itself and other objects. Again, properties of interacting objects and the liquid itself are likely to affect the sounds. For instance, the degree of splashing in a current may indicate the roughness or complexity of its enclosing channel, or the amount of wind over its surface. The liquid's viscosity also produces audible effects on the sounds it makes: it seems easy to tell whether a liquid gurgling out of a bottle is water, or a thicker syrup or oil. Such attributes probably affect high-level characteristics of the sounds' temporal evolution (e.g., the speed with which the frequencies and amplitudes change), while the style of this evolution itself may tie all liquid sounds together.

Though the physical attributes of aerodynamic and liquid sounds are not the same as those of vibrating solids, they do share common features. As with solids, these attributes can be grouped according to whether they belong to the causal interaction, the nature of the material involved, or the configuration of the material. Causal interactions are primarily responsible for the nature of aerodynamic sounds such as explosions, hisses and fan noises. Liquid sounds also depend on properties of their causal interactions, such as the force with which an object falls into liquid or the rate with which a liquid is poured. Material attributes are also important for both liquid and aerodynamic sounds, though viscosity differences among liquids are probably more commonly experienced than density differences among gasses. Finally, the configuration of the gas or liquid also has salient effects on the sounds they produce: drips or splashes on a shallow layer of liquid are likely to involve relatively high-frequency sounds, while the shape of a body of gas determines its modal vibration patterns and thus its resonating frequencies.

Complex Sound-Producing Events

Although all sound-producing events seem to involve vibrating solids, aerodynamic, or liquid interactions, many also depend on complex patterns of the simple events described above. So footsteps consist of temporal patterns of impact sounds, while door slams involve the squeak of scraping hinges and the impact of the door on its frame. Other sound-producing events involve more than one sort of material, as when liquid falls upon some solid surface. Though the discussion above may point to a useful framework for understanding the attributes of simple sounds, it does not directly address those of more complex sound-producing events.

Traditional physical accounts of sound-producing events do not describe these sorts of complex events, but there are higher-level physical attributes of such events that make reliable effects on their sounds. Some of these involve timing of successive events, so that, for instance, successive footstep sounds probably must occur within a range of rates and regularities to be heard as walking. Others are likely to involve mutual constraints on the objects that participate in related events. For instance, concatenating the creak of a heavy door closing slowly with the slap of a light door slammed shut would be likely to sound quite unnatural. These sorts of higher-level attributes of the events are not the sorts of variables that physicists typically study, but they are important for understanding everyday listening.

As with all events, sound sources may be considered nested in terms of scale (Warren & Shaw, 1985). We can listen to the sound created by an approaching automobile, its engine, or perhaps one of its cylinders (especially salient when something has gone wrong). But the attributes of sound-producing events change with scale. More complex events involve more complex and constrained combinations of interactions and materials. From this perspective, a *basic level* sound-producing event can be defined as one involving a single interaction and a single sound-producing object. Examples of basic level events might include hitting a solid, scraping it, explosions, and dripping noises. More complex events, then, can be understood in terms of combinations of basic level ones, combinations which are structured in ways which add information to their simpler constituents. From this point of view, it should be clear that complexity here is defined in terms of the event, not the ease of its perception: complex sounds are likely to be richer in structure than simple ones and thus easier to obtain information from (Jenkins, 1985). I will return to the subject of complex sound-producing events later in this paper.

WHAT DO WE HEAR?

Considering the physical attributes of sound-producing events is useful in driving intuitions about the sorts of perceptual dimensions and features that might characterise everyday listening. But knowing how the physics of an event determines the sound it makes is not the same as knowing how a sound specifies an event. For instance, several attributes of a vibrating solid, including its size, shape, and density, determines the frequencies of sound it produces. When hearing two sounds composed of different frequencies, then, how are we to know which parameters of the source has changed?

There are several possible answers to this problem. First, when several “basic” physical properties of an event have the same effect on the sound it produces, it is possible that a single perceptual dimension is heard, one which incorporates all of them. In other words, we may not hear size, shape, or density separately, but rather a new dimension (perhaps moment of inertia) which combines all of them. On the other hand, it may be that the effects of changing some attributes are much smaller than changing others – for instance, typical variations in object density are likely to make smaller changes in frequency than average changes in either size or shape. In this case, the parameters which cause larger variations may also be more salient. That is, we may be inclined to hear a change in frequency as a change in size rather than density, just because size changes are the more significant source of frequency changes. Finally, it is quite likely that many parameters that change frequency also change other attributes of a sound. For example, changing the size of an object will change the frequencies of the sound it produces, but not their pattern. Changing the shape, on the other hand, changes both the frequencies and their relationships. These complex patterns of change may serve as information distinguishing the physical parameters responsible: These are the acoustic invariants that an ecological acoustics should be concerned with discovering.

In any case, we can not base an account of everyday listening on the physics of sound-producing events alone. As Gibson (1966) pointed out, what is simple for physics may not be simple for perception, and vice versa. Instead, it is necessary to build an “ecological” physics, one founded on attributes relevant to listeners. For this reason, several studies have aimed at exploring the kinds of information sound conveys. Experiments of this sort complement analyses of physics. The data from such empirical studies constrain the sorts of physical attributes we might think we hear, while physical analyses can help in interpreting and organising experimental data.

Asking People What They Hear

One approach to understanding the information people hear is basically experiential, involving introspection and self-observation. For instance, Jenkins (1985) reported that blindfolded students were able to orient themselves within their environment on the basis of auditory information such as "acoustic landmarks," resonances and echoes, and mixtures of near and far sounds. Note that these students were not only using their ability to localise and to use reverberation as information about the environment (as discussed above), but that the sounds themselves – people talking, relatively continuous machine noises and the like – served as meaningful and relatively stable landmarks.

Another approach to understanding what people hear is simply to ask them. For instance, Vanderveer (1979) presented people with recorded tokens of 30 everyday sounds such as clapping and tearing paper in a free identification task. Participants in the study were run in groups, and asked to write a short phrase describing each sound. She found that they tended to identify the sounds in terms of the objects and events which caused them, describing their sensory qualities only when they could not identify the source events. In addition, people's mistakes tended to be based on the temporal qualities of sounds, so that clapping might be confused with dropping a book, but seldom with tearing paper.

I ran a similar study (Gaver, 1988) in which 17 sounds were played to people asked to describe what they heard. In contrast to Vanderveer's study, participants were run individually, and prompted by the experimenter to go into as much detail as they could about their experience. Like Vanderveer, I found that people nearly always described the sounds in terms of their sources. Moreover, their accuracy was often impressive. For instance, several participants could readily distinguish the sounds made by running upstairs from those made by running downstairs; others were substantially correct about the size of objects dropped into water; and most could tell from the sound of pouring liquid that a cup was being filled. Participants did find that some sounds were extremely difficult to identify (e.g., the sound of a file drawer being opened and closed), but they were almost always correct about some others (e.g., the sound of writing with chalk on a chalkboard). When mistakes were made, they often revealed interesting attributes that were heard. For instance, several people said the file drawer sounded like a bowling alley, both of which might be described as “rolling followed by impact(s)." Attempts to identify unusual or implausible sounds were equally interesting. For example, the sound of somebody walking across a floor covered with newspaper was described variously as a person walking on snow or gravel or as somebody rhythmically crumpling paper; only one participant identified the sound correctly, and then immediately rejected the correct perception as being too implausible. Results such as these are valuable in suggesting more focused empirical investigations and more generally in indicating that peoples' judgements correspond well to the physical account of events described above.

Ballas and his colleagues (Ballas & Howard, 1987; Ballas & Sliwinski, 1986) have used free identification tasks to study everyday sounds that are ambiguous as to their sources. They have shown that a measure of the information inherent in a given sound, based on the number of possible sources that people propose, can be used to predict reaction times for its identification.

Studies such as these are informative, but sometimes frustrating. The result of asking people what they hear is often a list of events or attributes more akin to a list of sound-effects (see below) than to a generative set of dimensions and attributes. It seems clear, given the difficulty in analysing such data, that future protocol studies should be done in conjunction with analyses of the physics of sound-producing events. Physical analyses are useful both in generating hypotheses about what source attributes are audible and in constraining the range of sounds used in protocol studies – or other empirical investigations – to focus on attributes of interest.

Nonetheless, asking people what they hear is invaluable as a complement to more analytic approaches. In many of these studies inherent perceptual categories reveal themselves both in correct answers and (perhaps even more often) in confusions. So for instance, the fact that the file drawer sound used in Gaver (1988) was confused with that of bowling suggests that rolling may be a particularly salient event warranting further examination. Such results, taken in conjunction with those of physical analyses, provide the foundation for a preliminary framework that describes the audible attributes of sound-producing events.

Describing Sound-Producing Events

In trying to understand what we hear in the world, the purpose is to characterise a wide range of the information we obtain by listening to sources. My concern is to organise the attributes of everyday listening in a form that is relatively simple and general. Just as we can capture a great deal of the sensory qualities of sounds with descriptions of their pitch, duration, loudness and so forth, so we would like to find a set of descriptive dimensions and features that characterise everyday listening. And just as qualities such as pitch and loudness apply generally to most musical sounds, we would like our dimensions and features to apply to broad ranges of everyday sounds.

What sort of descriptions will do? The vast range of everyday sounds make simple descriptions of them difficult. Consider Figure 5, for instance, which shows an excerpt from the table of contents of a sound-effects CD (Delta Music, 1987). This is a domain in which descriptions of everyday sounds has developed out of necessity. The first thing to notice is that the number of distinctive sounds listed is quite large – over 50 in this example, over 100 on the disk, and this disk is only one of many. The world of everyday sounds is immense.

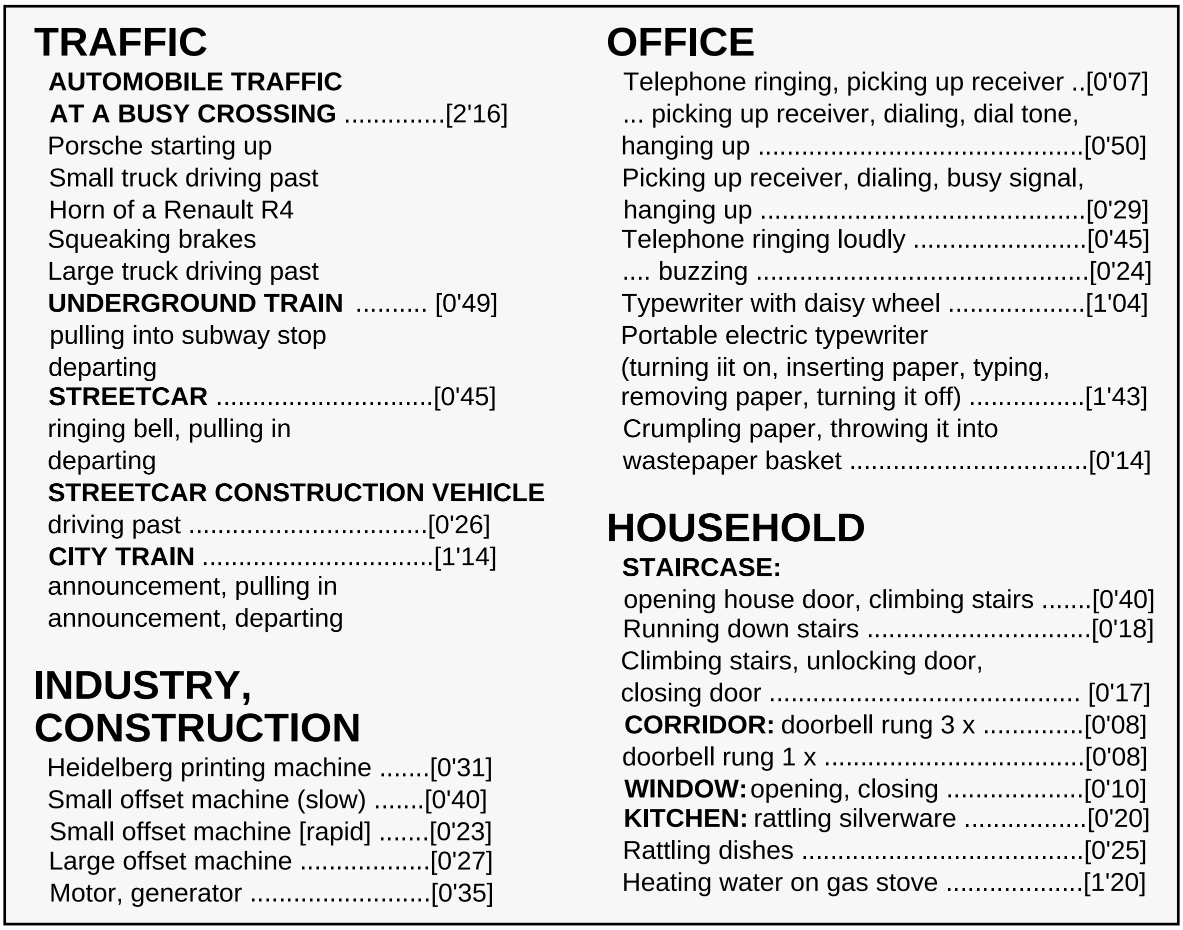

TRAFFIC

AUTOMOBILE TRAFFIC
AT A BUSY CROSSING ....[2'16]
Porsche starting up
Small truck driving past
Horn of a Renault R4
Squeaking brakes
Large truck driving past
UNDERGROUND TRAIN .... [0'49]
pulling into subway stop
departing
STREETCAR ....[0'45]
ringing bell, pulling in
departing
STREETCAR CONSTRUCTION VEHICLE
driving past ....[0'26]
CITY TRAIN ....[1'14]
announcement, pulling in
announcement, departing

INDUSTRY, CONSTRUCTION

Heidelberg printing machine ....[0'31]
Small offset machine (slow) ....[0'40]
Small offset machine [rapid] ....[0'23]
Large offset machine ....[0'27]
Motor, generator ....[0'35]

OFFICE

Telephone ringing, picking up receiver ..[0'07]
... picking up receiver, dialing, dial tone,
hanging up ....[0'50]
Picking up receiver, dialing, busy signal,
hanging up ....[0'29]
Telephone ringing loudly ....[0'45]
.... buzzing ....[0'24]
Typewriter with daisy wheel ....[1'04]
Portable electric typewriter
(turning iit on, inserting paper, typing,
removing paper, turning it off) ....[1'43]
Crumpling paper, throwing it into
wastepaper basket ....[0'14]

HOUSEHOLD

STAIRCASE:
opening house door, climbing stairs ....[0'40]
Running down stairs ....[0'18]
Climbing stairs, unlocking door,
closing door .... [0'17]
CORRIDOR: doorbell rung 3 x ....[0'08]
doorbell rung 1 x ....[0'08]
WINDOW: opening, closing ....[0'10]
KITCHEN: rattling silverware ....[0'20]
Rattling dishes ....[0'25]
Heating water on gas stove ....[1'20]

Figure 5. Excerpts from a sound-effects CD. Notice the mixed contextual, hierarchical, and dimension-based organisation.

But notice that the sounds in this example are organised to some extent. Most are listed in terms of the *context* in which they are likely to be heard: in traffic, in an industrial setting, etc. This may be useful for finding a desired sound, but it seems a poor basis on which to build a description of what we hear. For instance, the categories are not mutually exclusive; it is easy to imagine hearing the same event (e.g., a telephone ringing) in an office and a kitchen. Nor do the category names constrain the kinds of sounds very much. We might expect to hear anything from running water to a small appliance in a kitchen, with the only unifying feature being their supposedly typical environment. Domain-based descriptions of sound-producing events seem unlikely to provide an adequate description of the attributes of everyday listening.

More interesting are the suggestions of a hierarchical description of sounds. For instance, “automobiles” might be a superordinate category, with “Porsche” and “Renault R4” as subordinates. The list also suggests some dimensions (e.g., small and large offset machines) and features (e.g., printing machine, offset machine). A framework based on these sort of entities – hierarchies, features, and dimensions – seems a more promising approach than one based on context. Superordinate categories based on types of *events* (as opposed to contexts) provide useful clues about the sorts of sounds that might be subordinate, while features and dimensions are a useful way of describing the differences among members of a particular category. A hierarchical framework that describes sounds' attributes and dimensions thus seems more likely to be generative, to delineate a *space* of possible sounds, than context-based classifications.

Understanding the physics of sound-producing events is useful in suggesting physical attributes that might be heard, while the experiential and protocol studies described above help to constrain hypotheses about the attributes that are actually heard. Using our knowledge of physics and the results of these studies together, then, we may begin to build up a map of everyday sounds, a framework for understanding some of the basic source parameters they convey. Such a framework is necessarily speculative and probably incomplete. Nonetheless, it is useful in providing an organising account of the vast range of sound-producing events, in suggesting what source attributes might be revealed by the sounds they make, and in providing hypotheses for future research.

A Map Of Everyday Sound-Producing Events

To begin mapping the space of everyday sounds, sound-producing events are distinguished first by broad classes of materials and then by the interactions that can cause them to sound (Figure 6). Most generally, sounds indicate that something has happened, that an event has occurred, that there has been an interaction of materials. All sounds, then, convey this information.

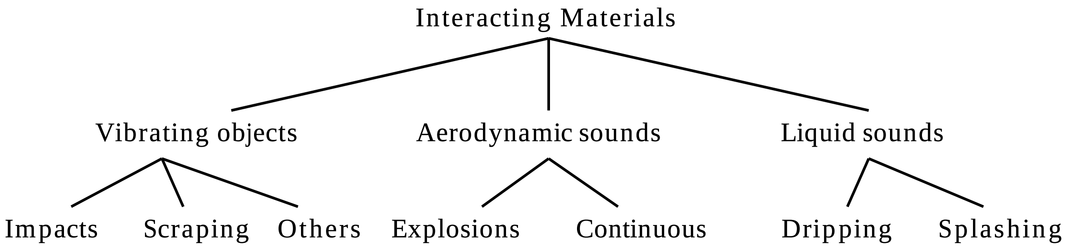

Figure 6. A hierarchical description of simple sonic events.

At the next level, sound-producing events may be divided into three general categories: those involving vibrating solids, aerodynamics, and liquids. This categorisation is supported both by the account of physics and the results of the protocol studies described above. Although participants often misinterpreted the sources of sounds they heard, no mistakes crossed these categories: Nobody confused the sounds made by vibrating solids, for instance, with those made by water. (This is not to say that such confusions never occur. In Mexico, for instance, "rain sticks" are constructed by inserting rows of wooden pegs inside a tube. When the tube is turned over, small beads and shells run down its length, striking these pegs and producing a sound remarkably like that of running water. This is an example of an illusion in everyday listening of the sort exploited by Foley artists creating sound-effects; see Mott, 1990.)

Finally, basic level sound-producing events are shown at the third level of this hierarchy, defined by the simple interactions that can cause solids, gasses and liquids to sound. For instance, the sounds made by vibrating solids may be caused by impacts, scraping, or other interactions. Aerodynamic sounds may be made by discrete, sudden changes of pressure (explosions), by more continuous introductions of pressure variations (gusts and wind). Similarly, liquid sounds may be caused by discrete drips, or by more continuous splashing, rippling, or pouring events.

This simple classification is the basis of the more comprehensive map of everyday sounds shown in Figure 7. This figure is broken into three main overlapping regions, corresponding to sound-producing events involving solids, liquids, and gasses respectively. Within each region, complexity grows towards the centre, with basic level events shown on the periphery and three kinds of complexity, described below, shown at progressively more central locations. The overlapping regions in the centre of the figure show examples of sound events that involve hybrids of different sources.

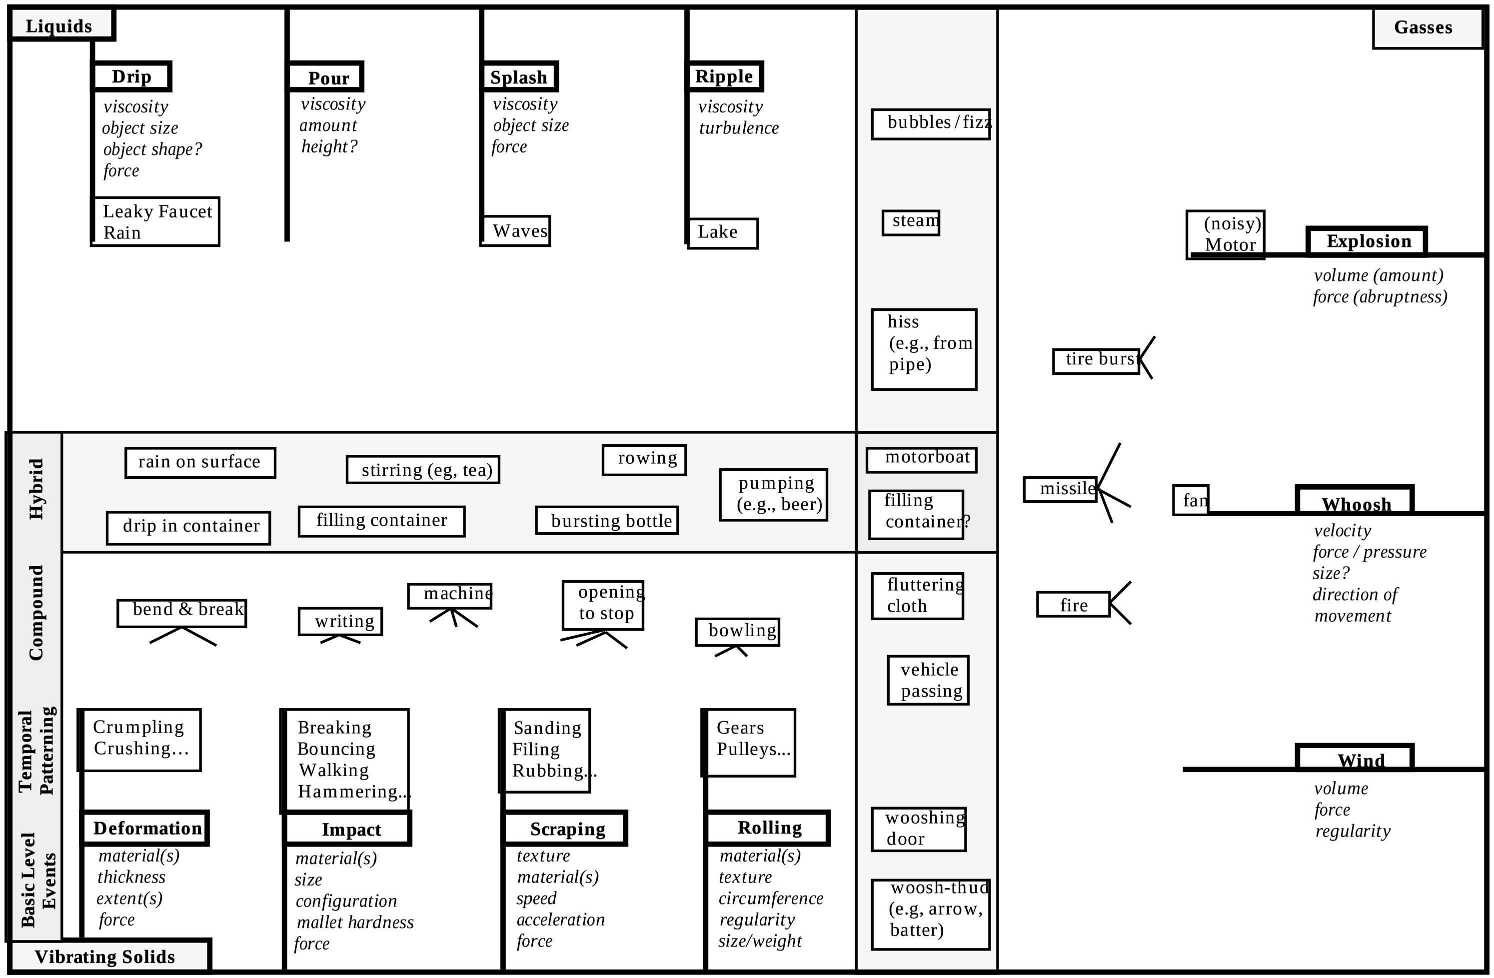

Figure 7. A "map" of everyday sounds. Three fundamental sources (vibrating solids, liquids and aerodynamics) are shown in the three overlapping sections of the figure. Within each section, basic sound-producing events are shown in bold, and their relevant attributes next to them in italics. Complexity grows towards the centre of the figure, with examples showing temporally patterned, compound, and hybrid sounds.

Basic Level Sources

Consider, for example, the region describing sounds made by vibrating solids. Four fundamentally different sources of vibration in solids are indicated as basic level events: deformation, impacts, scraping and rolling. Under each of these basic level events are listed the attributes of these events that seem relevant for the sounds they produce. For instance, impact sounds seem to convey information about five aspects of the event: the vibrating solid's material, size and configuration, the surface hardness of the impacting materials, and the force of the impact. Scraping, on the other hand, conveys information about the texture of the surfaces, one or more of the materials involved, the speed and acceleration of scraping, and the force with which one object is scraped over another.

Patterned, Compound, and Hybrid Sources

Above these basic level events three sorts of complex events are shown. The first is defined in terms of *temporal patterning* of basic events. For instance, breaking, spilling, walking and hammering are all complex events involving patterns of simpler impacts. Similarly, crumpling or crushing are examples of patterned deformation sounds. We would expect these sorts of events to convey the attributes made by their basic level constituents; in addition, other sorts of information are made available by their temporal complexity. For example, the regularity of a bouncing sound provides information about the symmetry of the bouncing object, while variations in the scraping sounds produced by filing might indicate the general configuration of the object being filed.

The next level of complexity is produced by *compound* events which involve more than one sort of basic level event. For instance, the sounds made by writing involve a complex series of impacts and scrapes over time, while those made by bowling involve rolling followed by impact sounds. Again, these sounds are likely to convey information inherited from their basic level components as well as new information made available by their complexity. It is also worth noting that while some of these events involve more than one sort of source simultaneously (e.g., writing), others involve a series of basic events (e.g., opening a file drawer until in impacts against its stops).

Finally, *hybrid* events involve yet another level of complexity in which more than one basic sort of material is involved. For instance, when water drips on a reverberant surface, the resulting sounds are caused both by the surface's vibrations and the quickly-changing reverberant cavities, and thus involve attributes both of liquid and vibrating solid sounds. Some hybrid events involve attributes of all three basic sources: for instance, the sounds made by a speeding motorboat involves the splashing of the water, the vibrating engine sounds, as well as the rush of air past the body of the boat. As with other complex events, hybrid events provide information about their materials and the basic level events involved in their production, as well as more idiosyncratic information specific to their sources.

Developing the Framework

Clearly the framework shown in Figure 7 is far from complete. For one thing, we know much more about how to characterise the sounds produced by vibrating solids than we do about either liquid or aerodynamic sounds. In part this is a problem of language (what do you call the continuous rippling sound heard on a lake?) In addition, we don't yet know how to characterise the attributes of many complex events (though see Warren & Verbrugge, 1984). Finally, the three basic categories of sounds shown here may not be enough. What of electronic sounds, such as those made by sparks or humming wires? Should fire be a basic sound-producing category (as suggested by the Aristotelian earth-air-water trichotomy used here)? What of vocal sounds? Though the framework proposed here seems powerful and correct, there is much more to be discovered about how people hear events.

There are several more fundamental problems with this system. Two have to do with the sources of information used in its creation. Insofar as this map relies on verbal evidence from subjects, it is liable to confuse the effects of language use for the attributes of perceptual experience. Insofar as it is based on an analysis of physics, it is liable to confuse source attributes that affect sounds with those that are actually heard. For example, the texture of an object that falls into a liquid certainly affects the physics of dripping, but it is by no means obvious that this property is perceptible. Clearly, more focused empirical examinations are needed to bring the descriptions offered by everyday language and analytical physics closer together.

In addition, the idea that the attributes of different sorts of events may be catalogued unequivocally is somewhat questionable. The information people obtained from the various sound-producing events used in my protocol study often seemed to be somewhat peculiar to the particular sounds used as examples. Not all impact sounds, for instance, provided equal information about material, while some sounds conveyed more information about materials than interactions (for instance, many subjects heard metal but not deformation when listening to a crumpling can). For this reason, information for the source attributes shown under the various basic level events in this framework should perhaps be thought of as possibly, rather than necessarily, available in particular examples of the sounds.

Nonetheless, this framework does describe a large range of sound-producing events. It does so in a way that recognises the mutual constraints of materials and interactions in producing basic level events, and shows some of the attributes we might expect to hear from them. In addition, it provides a mechanism for understanding the myriad of complex sound-producing events, and through its inheritance mechanism provides a succinct account of some of the attributes such sounds might make available. Finally, and perhaps most important, it offers a way of understanding the relations among existing results concerning everyday listening, and suggests new avenues for exploration.

CONCLUSIONS

I hope to have made clear in this paper that everyday listening and ecological acoustics is a vast, rich, and relatively unexplored territory. Moreover, I have tried to sketch out a crude map for its understanding. In discussing the continuum between event and audition, my concern has been to evoke intuitions about the structured information provided by sound about events at locations in an environment, and to point out the relatively direct route from the structuring event to the auditory system. In presenting more detailed, if qualitative, accounts of the physics of sound-producing events and the results of protocol studies, my aim has been to explore further the particular attributes of sources that we might hear. Finally, in presenting a map of the range of sound-producing events, I hope to have summarised this domain in a form that will be useful for its further exploration.

Like most early maps of relatively unexplored domains, the one offered here will undoubtedly be revised as we learn more about everyday listening and ecological acoustics. It is likely to be incomplete and wrong at various levels of generality. Even the basic taxonomy on which it is based may be altered if a new system is found. But this is not only to be expected, it is to be hoped for; the purpose of maps such as these is to entice other explorers into the domain.

For if it is clear that there is much more to be learned about everyday listening, it is equally clear that such knowledge is valuable. Already basic research in this area has found application in the form of *auditory icons*, everyday sounds used to convey information about computer events by analogy with everyday sound-producing events (see, e.g., Gaver 1991, 1986). For instance, selecting a file icon might make the sound of tapping an object, with the struck object's size indicating the file size and the object's material indicating the type of file. Auditory icons such as this use the dimensions of everyday listening in mapping sounds to events in computer systems, and thus can provide closer, more intuitive mappings than strategies which rely on musical dimensions. In addition, because the sounds used for auditory icons are caricatures of those we might hear in the everyday world, they tend to fit better into the audio ambience in which such systems may be used. Research on everyday listening and auditory icons complement one another, with basic research providing information about what we hear and the acoustic information we use, while applications suggest the attributes and dimensions of everyday listening which are likely to be particularly relevant. But it is important to note that the notion of everyday listening came first, with auditory icons following as a natural application of these ideas.

Research on the psychology of everyday listening is valuable in its own right as well, balancing the typical bias towards studying vision in understanding how people perceive and act in the world. Clearly, there is much to be done. In the beginning of this paper, I suggested that an understanding of everyday listening should be guided by two questions: what do we hear? and how do we hear it? My intention here has been to address the first of these questions, with the map of sound-producing events presented here the result of this exploration, one that will undoubtedly be refined and expanded with further study. But in any case, this is only half the story: We also need to know how we hear the things we do, describing the acoustic information that is available for them. Research reported by Warren & Verbrugge (1984), Freed & Martens (1986), Wildes & Richards (1988), and Shaw et al. (1991) focus on this topic, and it is the subject of a companion paper to this one (Gaver, in preparation). Nonetheless, we know relatively little about a relatively small set of source events. This should be no surprise: Everyday listening and ecological acoustics remain a relatively untrammelled theoretical and empirical domain awaiting further exploration.

ACKNOWLEDGEMENTS

This work has developed over several years, with parts of this paper being based on my doctoral dissertation (Gaver, 1988). I am indebted to many people for their support, comments and insights. In particular, I thank Don Norman, Anne Schlottmann, Wendy Mackay and Alan MacLean. Finally, I am grateful to Bill Mace and Michael Turvey for their enthusiasm and support of this research.

REFERENCES

Ballas, J. A., & Howard, J. H., Jr. (1987). Interpreting the language of environmental sounds. Environment and Behavior. 19, 91-114.

Ballas, J. A., & Sliwinski, M. J. (1986). Causal uncertainty in the identification of environmental sounds. (Tech. Rep. ONR-86-1). Washington, DC: Georgetown University, Department of Psychology.

Blauert, J. (1983). *Spatial hearing: The psychophysics of human sound localization*. Cambridge, MA: MIT Press.

Cage, J. (1976). *Silence: Lectures and writings by John Cage*. Middletown, CT: Wesleyan University Press.

Cremer, L. (1984). *The physics of the violin.* Trans. by John S. Allen. Cambridge, MA: The MIT Press.

Delta Music (1987). *Digital sound tracks: 80 Geräusche und Klangeffeckte*. Vol. 1.

Edgerton, H. & Killian, J. Jr. (1975). *Moments of vision.* Cambridge, MA: The MIT Press.

Evans, E. F. (1982). Functional anatomy of the auditory system. In H. B. Barlow and J. D. Mollon (Eds.), *The Senses.* Cambridge, England: Cambridge University Press.

Freed, D. J., & Martens, W. L. (1986). Deriving psychophysical relations for timbre. *Proceedings of the International Computer Music Conference*. Oct. 20 - 24, 1986, The Hague, The Netherlands.

Gaver, W. (in preparation). How do we hear in the world? Acoustic information for source attributes. To be submitted to *Ecological Psychology.*

Gaver, W. W. (1986). Auditory icons: Using sound in computer interfaces. *Human-Computer Interaction.* 2, 167 - 177.

Gaver, W. W. (1988). Everyday listening and auditory icons. Doctoral Dissertation, University of California, San Diego.

Gaver, W. W. (1991). Sound support for collaboration. *Proceedings of the Second European Conference on Computer-Supported Collaborative Work* (Amsterdam, September 24 - 27, 1991) Kuwer, Dordrecht.

Gibson, J. J. (1966). *The senses considered as perceptual systems.* Boston: Houghton Mifflin.

Gibson, J. J. (1986). *The ecological approach to visual perception.* Hillsdale, New Jersey: Lawrence Erlbaum Associates (original version published in 1979).

Guski, R. (1990). Auditory localization: effects of reflecting surfaces. *Perception*, *19*, 819 - 830.

Hafter, Ervin R., & Buell, Thomas N. (1985). The importance of transients for maintaining the separation of signals in auditory space. Posner, Michael I., & Marin, Oscar S. M. (Eds.), *Attention and Performance XI.* Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Handel, S. (1989). *Listening: An introduction to the perception of auditory events.* Cambridge, MA: The MIT Press.

Heine, W-D, & Guski, R. (1991b). Preliminaries to a concept of ecological acoustics. In Beek, P., Bootsma, R., and van Wieringen, P., (Eds.): *Studies in perception and action: Posters presented at the VIth international conference on event perception and action.* Amsterdam: Rodopi.

Heine, W-D., & Guski, R. (1991a). Listening, the perception of auditory events? An essay review of Listening: An Introduction to the Perception of Auditory Events by Stephen Handel. *Ecological Psychology. 3*(3), 263 - 275.

Helmholtz, H. L. F. (1885/1954). *On the sensations of tone as a physiological basis for the theory of music.* New York: Dover.

Hoeger, R. (1992). The role of acoustical texture in perceiving auditory events. Paper presented at the Second European Conference On Ecological Psychology (University of Strathclyde, Glasgow, 24 - 26 June 1992).

Jenkins, James J. (1985). Acoustic information for objects, places, and events. In Warren, W. H., & Shaw, R. E., *Persistence and change: Proceedings of the first international conference on event perception.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Lehnert, H., & Blauert, J. (1991). Virtual auditory environment. In the proceedings of the IEEE workshop.

Lindsay, P. H., & Norman, D. A. (1977). *Human information processing.* (2nd ed.). New York: Academic Press.

Mott, R. L. (1990). *Sound effects: Radio, TV, and film.* London: Focal Press.

Pierce, John R. (1983). *The science of musical sounds*. New York: W. H. Freeman and Company.

Rayleigh, Lord (1877/1929). *Theory of sound.* London: Macmillan

Runeson, S. (1977). On the possibility of "smart" perceptual mechanisms. *Scandinavian Journal of Psychology.* 18, 172-179.

Shaw, B., McGowan, R., & Turvey, M. (1991). An acoustic variable specifying time-to-contact. *Ecological Psychology, 3*(3), 253 - 261.

Skudrzyk, E. (1968). *Simple and complex vibratory systems.* University Park, PA.: The Pennsylvania State University Press.

Van Dyke, M. (1982). *An album of fluid motion.* Stanford, CA: The Parabolic Press.

Vanderveer, N. J. (1979). Ecological acoustics: Human perception of environmental sounds. Dissertation Abstracts International. 40/09B, 4543. (University Microfilms No. 8004002).

Warren, W. H., & Shaw, R. E. (1985). Events and encounters as units of analysis for ecological psychology. In Warren, W. H., & Shaw, R. E., *Persistence and change: Proceedings of the first international conference on event perception.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Warren, W. H., & Verbrugge, R. R. (1984). Auditory perception of breaking and bouncing events: A case study in ecological acoustics. *Journal of Experimental Psychology: Human Perception and Performance.* 10, 704 - 712.

Wenzel, E., Wightman, F., & Kistler, D. (1991). Localization with non-individualized virtual acoustic display cues. In proceedings of CHI, 1991 (New Orleans, Louisiana, April 28 - May 2, 1991) ACM, New York, 351 - 359.

Wildes, R., & Richards, W. (1988). Recovering material properties from sound. Richards, W. (ed.), *Natural computation*. Cambridge, MA: MIT Press.