Chapter 1

Everyday listening: an annotated bibliography

Bruno L. Giordano
Università di Udine – Faculty of Education
Udine, Italy
Università di Padova – Department of General Psychology
Padova, Italy
bruno.giordano@unipd.it

There is an immense world of sounds that have been scarcely considered by traditional research in auditory perception, the world of everyday sounds. We should systematize it, understand its own regularities and laws.

Many of the sounds we encounter every day are generated by physical events that involve an interaction between objects (a coin dropped on the floor, water drips falling on the sink) or a change of the properties of single objects (a bursting balloon). Are listeners able to recover the properties of these physical events on the basis of auditory information alone? What is the nature of the information used to recover these features? These questions were originally raised inside the ecological approach to , firstly extended to the study of the auditory modality by Vanderveer [240]. Gaver [92, 97] outlined a taxonomy of environmental sounds, in order to "entice other explorers" into this domain. His taxonomy is based on the simple assertion that sounds are generated by an "interaction of materials", and it is inside this map that we should place the pieces of experimental evidence collected so far, synthesized in this chapter.

Research in sound source recognition requires methodological tools, of course. A formalization of the research design in this field is found in [153]. The object under study can be described at three levels: the physical, acoustical and perceptual ones. All the pairwise investigations between these three levels of description should receive our attention. Analysis of the relationship between the perceptual and physical levels tells us if the investigated feature of the physical event is properly recognized or scaled by listeners. Analysis of the relationship between the physical and acoustical levels tells us which properties of the acoustical signal differentiate between categories or levels of the physical feature of interest. Analysis of the relationship between the acoustical and perceptual levels tells us whether the acoustical information outlined at the previous stage is effectively used by participants to scale or categorize the physical feature of interest or whether other acoustical features influence listeners responses. This last analysis should be validated by manipulation of the acoustical properties found significant in explaining participants responses.

A corollary to Li et al. [153] formalization is that we should assess whether recognition of a particular feature of the sound source is possible despite variations in extraneous physical features. In other words we should make use of what we define as *perturbation variables*, defined as those variable features of the sound source, extraneous from the one whose recognition is investigated. For example Kunkler-Peck and Turvey [142] studied scaling of the dimensions of struck plates upon variations in their dimensions and material. In this case material is a perturbation variable.

Finally we would like to make the reader aware of one more issue. Research on sound source recognition could demonstrate several misalignments between the physical features of the sound source and the recognized ones, especially if perturbation variables are extensively used. The misalignment between physical and perceptual levels is a well known lesson in the visual perception field, and there will be no surprise to empirically demonstrate this in the auditory domain. In this case research on sound source recognition will turn as research on the *apparent features of the sound source*. In this case we will have to start thinking about this type of research as an alternative way for addressing a problem which is still widely unknown: the problem of the quality of auditory percepts. As the problem of quality in audition has been already addressed by research on timbre perception (see [171, 113] for a review), we would like to conclude this introduction using a metaphor in debt with these studies. The world of sounds can be conceived as organized within an multidimensional space. Classical research on auditory perception have investigated this space along dimensions such as pitch, loudness, duration, timbral brightness, attack hardness and so on. Research on sound source recognition envisions the existence of a set of new dimensions to investigate this multidimensional space, criteria based on the features of the sound source, be them coherent with the physical reality or apparent.

1.1 1979

S. J. Lederman. *Auditory texture perception* [148]

Lederman compared the effectiveness of tactile and auditory information in judging the roughness of a surface (i.e., the *texture*). Roughness of aluminum plates was manipulated by varying the distance between adjacent grooves of fixed width, or by varying the width of the grooves. The subjects' task was to rate the roughness of the surface numerically. In the auditory condition, participants were presented the sounds generated by the experimenter who moved his fingertips along the grooved plate. In the remaining conditions, subjects were asked to move their fingertips onto the plate: In the tactile condition they wore cotton plugs and earphones while touching the plate; in the auditory-plus-tactile condition they were able to hear the sound they generated when touching the plate.

Roughness estimates were not different between the auditory-plus-tactile and tactile conditions, but differed in the auditory condition. In other words when both kinds of information were present, the tactile one played the strongest role in determining experimental performance. Roughness estimates were shown to increase as both the distance between grooves and the width of the grooves decreased. Additionally roughness estimates increased as the force exerted by the finger over the surface increased, and the speed of the finger movement decreased. The effect of the force on roughness estimates in the auditory condition was however not constant across subjects. A speculative discussion concerning the relative role of pitch and loudness in determining the estimates is provided by the author, although no acoustical analyses of the experimental stimuli are provided.

N. J. Vanderveer. *Ecological Acoustics: Human perception of environmental sounds* [240]

Vanderveer's dissertation applies for the first time the ecological approach to auditory perception. Empirical evidence was collected using three different methodologies. In the first study a tape containing 30 environmental sounds (e.g., crumpling paper, jingling coins, wood sawing) was played continuously to two groups of subjects. They were asked to write down what they heard. Descriptions of the event were given by subjects based on the sound features, rather than the acoustical properties of the stimuli. However, subjects were found to refer to abstract perceptual qualities when they did not recognize the sound source. Accuracy was fairly good in all cases. Descriptions focused mainly on the actions causing the sound events, rather than the objects' properties. Confusions were frequent among stimuli sharing common temporal patterns, when such patterns were hypothesized to carry information about the actions.

In a second experiment subjects were asked to classify 20 environmental sounds into a number of different categories, based on perceived similarity. Again, clustering seemed to be based on the similarity of the temporal patterns. Results collected using the same classification procedure were confirmed by another experiment rating similarity, that was conducted using a subset of the stimuli adopted in the classification experiment.

1.2 1984

W. H. Warren and R. R. Verbrugge. ***Auditory perception of breaking and bouncing events: a case study in ecological acoustics*** **[251]**

According to the terminology adopted by those embracing the ecological approach to perception, we distinguish between two classes of invariants (i.e., higher-order acoustical properties) that specify the sound generation event. Structural invariants specify the objects properties, whereas transformational invariants specify the way they change. Warren and Verbrugge investigated the nature of the structural invariants that allow identification of breaking and bouncing events. By conducting a physical analysis of these two kinds of events, the authors hypothesized that the nature of these invariants was essentially temporal, the static spectral properties having no role in the identification of breaking and bouncing events. Experimental stimuli were generated by dropping on the floor one of three different glass objects from different heights, so that for each object a bouncing event and a breaking one was recorded. Once the ability of the participants to correctly identify these two types of events was assessed using the original stimuli, two further experiments were conducted making use of synthetic stimuli. The bouncing event was synthesized by superimposing four damped, quasi-periodic pulse trains, each one generated by recording one of four different bouncing glass tokens. These four sequences exhibited the same damping coefficient. The breaking event was synthesized by superimposing the same pulse trains, this time using a different damping coefficient for each of them (in the second experiment the breaking stimuli were preceded by a 50 ms noise burst of the original breaking sound). The identification performance was extremely accurate in all cases, despite the strong simplifications of the spectral and temporal profile of the acoustical signal. Therefore the transformational invariants for bouncing were identified to be a single damped quasi-periodic sequence of pulses, whereas those for breaking were identified to be a multiple damped, quasi-periodic sequence of pulses.

1.3 1987

B. H. Repp. *The sound of two hands clapping: an exploratory study* [200]

Speech perception has been classified by Liberman and Mattingly as perception of phonetic gestures [154]. This theory has been given the name of motor theory of speech perception. Repp's work extends this theoretical approach to the investigation of non-speech communicative sound: claps. In particular, Repp hypothesized the subjects' ability to recognize the size and the configuration of clapping hands by auditory information. The recognition of the hand size was also put in relation with the gender recognition of the clapper, given that male have in general bigger hands than females. Several clapping sounds were recorded from different clappers. In the first experiment the recognition of the clapper' gender and his/her hand size was investigated indirectly, as participants were asked to recognize the clapper's identity. Recognition was not good, although the listeners' performance in the identification of their own claps was much better. Gender recognition was barely above chance. Gender identification appeared to be guided by misconceptions: faster, higher-pitched and fainter claps were judged to be produced by females and vice-versa. In the second experiment, subjects had to recognize the configuration of the clapping hands. Subjects were found to be able to recover correctly the hand configuration from sound. Although the hand configuration was significant in determining the clapping sound spectrum, nevertheless the best predictor of performance was found to be the clapping rate, the spectral variables having only a secondary role during the recognition task.

W. H. Warren, E. E. Kim, and R. Husney. ***The way the ball bounces: visual and auditory perception of elasticity and control of the bounce pass*** **[250]**

Warren et al. studied the perception of elasticity in bouncing balls, in both the visual and auditory domains. In experiment 1 subjects were requested to bounce a ball off the floor to a constant target height. Five balls with different elasticities were used, and subjects were exposed to different kinds of information concerning the elasticity of the ball before executing the task (in the auditory condition they heard the sound of a ball dropped on the floor, whereas in the auditory plus visual condition they saw and heard the same ball bouncing on the floor). Results showed that prior exposure to both visual and auditory information about the ball's elasticity did not lead to a different performance level, compared to when only auditory information was provided. In experiment 5 subjects had to rate the elasticity of a bouncing ball, simulated with a white circle bouncing on a white line at the bottom of the screen, and with a 40 ms, 190 Hz tone that was presented at visual impact time to simulate the bouncing sound. The rated bounciness was the same when it was judged from auditory or visual information only. In all cases, subjects appeared to judge elasticity by evaluating a single inter-bouncing period, rather than the ratio between successive periods.

1.4 1988

W. W. Gaver. ***Everyday listening and auditory icons*** **[92]**

In this section two empirical investigations reported by Gaver in his dissertation are summarized. In the first investigation participants were presented 17 different environmental sounds, generated by the interaction between solid objects and/or liquids. They were asked to describe what they heard, possibly providing the highest level of detail. Impacts were always identified: subjects were able to extract information concerning the object material, and possibly its size and hollowness. The accuracy in the identification depended on the specificity of the details. Crumpling can sounds were often confused with multiple impacts, probably because of the common physical nature of the two events. Interestingly, when crumpling sounds were generated using paper, confusion between crumpling and group of impacts was rare. Liquid sounds were correctly identified in all cases. Results gathered using more complex sounds again revealed a high performance level. For example, electric razor sounds were always recognized as being generated by a machine. Walking sounds were always correctly identified, with three subjects correctly recognizing sounds as being generated by a person walking upstairs, rather than downstairs.

Two experiments on real struck bar sounds were then performed. In the first experiment subjects had to categorize the material of struck metal and wooden bars of different lengths. Performances between 96 and 95% correct were observed. In the second experiment subjects had to estimate the length of the same bars from sound. A control group of subjects was, instead, asked to estimate the pitch of the same set of stimuli. Length ratings were almost a linear function of the physical length, the type of material being non-significant. For length estimates an interaction between material and length was found, so that the shortest and the longest metal bars were estimated to be shorter than wooden bars of equal length. The psychophysical functions for pitch scaling were much different for the two materials, the metal bar sounds having a higher pitch than the wooden ones. Similar results were found for all the tasks using synthetic bar sounds, except for the length rating. In this case, modeled changes of the bar length were associated to smaller variations in the estimated length compared to those observed using real sounds.

R. Wildes and W. Richards. ***Recovering material properties from sound*** **[253]**

The purpose of the authors was to find an acoustical parameter that could characterize material type uniquely, i.e. despite variations in objects features such as size or shape. Materials can be characterized using the coefficient of internal friction $tan\phi$, which is a measure of anelasticity (in ascending order of $tan\phi$ we have rubber, wood, glass, and steel). In the acoustical domain the coefficient of internal friction was found to be measurable using both the quality factor Q^{-1} and the decay time of vibration t_e, this latter measured as the time required for amplitude to decrease to $1/e$ of its initial value. For increasing $tan\phi$ we have an increase in Q^{-1}, and in t_e.

1.5 1990

D. J. Freed. ***Auditory correlates of perceived mallet hardness for a set of recorded percussive events*** **[82]**

Freed's study aims to measure an attack-related timbral dimension using a sound source-oriented judgment scale: hardness. Stimuli were generated by percussing four cooking pans, having variable diameter, with six mallets of variable hardness. Mallet hardness ratings were found to be independent of the pan size, thus revealing the subjects' ability to judge the properties of the percussor independently of the properties of the sounding object. The main goal of this study was to derive a psychophysical function by mallet hardness ratings, based on the properties of the acoustical signal. Preliminary experiments pointed out that significant information about mallet hardness was contained in the first 300 ms of the signals. For this reason, the acoustical analyses focused on this portion of the signals. Four acoustical indexes were measured: spectral level, spectral level slope (i.e., rate of change in spectral level, a measure of damping), average spectral centroid, and spectral centroid time weighted average (TWA). These acoustical indexes were used as predictors in a multiple regression analysis. Together, they accounted for 75% of the variance of the ratings.

1.6 1991

X. Li, R. J. Logan, and R. E. Pastore. *Perception of acoustic source characteristics: Walking sounds* [153]

Li et al. studied gender recognition in walking sounds. Walking sounds of seven females and seven males were recorded. Subjects were asked to categorize the gender by a four-step walking sequence. Results show recognition levels well above chance. Several anthropometric measures were collected on the walkers. Male and female walkers were found to differ in height, weight and shoe size. Spectral and duration analyses were performed on the recorded walking excerpts. Duration analysis indicated that female and male walkers differed with respect to the relative duration of the stance and swing phases, but not with respect to the walking speed. Nonetheless, judged masculinity was significantly correlated with the latter of these two variables, but not with the former. Several spectral measures were derived from the experimental stimuli: spectral centroid, skewness, kurtosis, spectral mode, average spectral level, and low and high spectral slopes. Two components where then derived by applying a principal components analysis on the spectral predictors. These components were used as predictors for both physical and judged gender. Male walkers in general were characterized by having a lower spectral centroid, mode and high frequency energy than females, and by higher values of skewness, kurtosis and low-frequency slope. The same tendencies were found when the two components were used as predictors for the judged gender.

Results gathered from the analysis of the relationship existing between the acoustical and the perceptual levels were then tested in another experiment. Stimuli were generated by manipulating the spectral mode of the two most ambiguous walking excerpts (also the spectral slopes were altered, but the manipulation of this feature was not completely independent of the manipulation of the spectral mode). Consistently with previous analyses, the probability of choosing the response "male" was found to decrease with increasing spectral mode. A final experiment showed that the judged gender could be altered by making a walker wear shoes of the opposite gender.

1.7 1993

W. W. Gaver. *What in the world do we hear? An ecological approach to auditory event perception* [97]

In the everyday life we recognize events and sound sources rather than sounds. This listening attitude has been defined by Gaver as "everyday listening", as opposed to "musical listening" where the perceptual attributes are those concerned by traditional research in audition. Despite the behavioral relevance of non-musical and non-speech sounds, empirical researches on them are missing. Research on everyday sounds focuses on the study of new perceptual features and dimensions, those concerning the sound source. Analyzing how the sound source features structure the acoustical signal is thus necessary, to find a set of candidate dimensions. This analysis, however, does not tell us which of these dimensions are relevant to everyday listening. For this purpose it is thus necessary to use protocol studies. The map of everyday sounds compiled by Gaver is based on both the knowledge about how a sound source structures the acoustical signal, as well as on protocol studies data. The most important distinction is found between solid, liquid and aerodynamic sounds, as protocol studies showed that these macro-classes are seldom confused each other. Then each of these classes is divided based on the type of interaction between materials. For example, sounds generated by vibrating solids are divided in rolling, scraping, impact and deformation sounds. These classes are "basic level sound-producing events". Each of them make different sound source properties evident.

The next level contains three types of complex events: those defined by a "temporal patterning" of basic events (e.g., bouncing is given by a specific temporal pattern of impacts); "compounds", given by the overlap of different basic level events; "hybrid events", given by the interaction between different types of basic materials (i.e., solids, liquids and gasses). Each of these complex events should potentially yield the same sound source properties, made available by the component basic events plus other properties (e.g., bouncing events may provide us informations concerning the symmetry of the bouncing object).

W. W. Gaver. ***How do we hear in the world? Explorations in ecological acoustics*** **[95]**

A study on everyday listening should investigate both the relevant perceptual dimensions of the sound generation events (i.e., what we hear) and the acoustical information through which we gather information about the events (i.e., how we hear). In timbre perception the study of relevant acoustical information can be based upon the so-called analysis and synthesis method [203]. This methodology looks for relevant information by progressively simplifying the acoustical structure of the signals investigated, until only the acoustical properties, whose further simplification would lead to relevant timbral changes, are retained. Likewise, everyday-sound synthesis algorithms are developed after the analysis of both the acoustical and physical event, complemented by the perceptual validation which has been made possible by the synthesis stage. Several algorithms are presented for the synthesis of impact, scraping, dripping, breaking/bouncing/spilling, and machine sounds. A final discussion highlights on some of the methodological issues that are connected to the validation of synthesis models.

1.8 1997

S. Lakatos, S. McAdams, and R. Caussé. ***The representation of auditory source characteristics: simple geometric form*** **[145]**

Lakatos et al. studied the listeners' ability to discriminate the shape of steel and wooden bars, as specified by the ratio between their height and width (H/W). Sounds were generated striking the bars with a mallet. All the stimuli were equalized in loudness. A cross-modal matching task was performed by the subjects, in which they had to indicate which one of two possible sequences of figures (representing the two bars with different H/W ratios) corresponded to the sequence heard. Stimuli generated by percussing steel and wooden bars were tested in different sessions. Subjects who did not reach a

75% correct criterion were excluded from further analyses (8.3% of the subjects did not reach that criterion for the steel bars, and 16.6% did not reach it for the wooden bars). The correct scores, converted in the appropriate way, were analyzed using MDS techniques. A two-dimensional solution was derived for the data related to steel bars. The coordinates labeling the first dimension were highly correlated with the H/W ratio, and with the frequency ratio of the transverse bending modes. The coordinates labeling the second dimension highly correlated with the spectral centroid. A cluster analysis of the same data set revealed a gross distinction between thick and thin bars (blocks vs. plates). Data gathered on wooden bars sounds led to a one-dimensional MDS solution, with the same correlation properties exhibited by the first dimension of the MDS solution derived in the case of steel bars.

R. A. Lutfi and E. L. Oh. *Auditory discrimination of material changes in a struck-clamped bar* [163]

Lutfi and Oh studied the discrimination of material in synthetic struck clamped bar sounds. Stimuli were synthesized by varying the parameters of bar elasticity and density toward characteristic values which characterize iron, silver, steel, copper, glass, crystal, quartz, and aluminum, respectively. Perturbations were applied either to all the frequency components (lawful covariation) or independently to each component (independent perturbation). On half of the trials participants had to tell which one of two presented stimuli was an iron sound, silver, steel, and copper being the alternatives. On the other half of the trials the target was glass, and the alternatives were crystal, quartz, and aluminum. Participants were given feedback on the correctness of their response after each trial. Performances were analyzed in terms of weights given to three different acoustical parameters: frequency, decay, and amplitude. The data revealed that discrimination was mainly based on frequency in all conditions, the amplitude and decay rate having only a secondary role.

1.9 1998

C. Carello, K. L. Anderson, and A. J. Kunkler-Peck. *Perception of object length by sound* [49]

Carello et al. investigated the recognition of the length of wood rods dropped on the floor. In two experiments, the former focusing on longer rods, subjects judged the perceived length by adjusting the distance of a visible surface put in front of them. Subjects were able to scale the rod length consistently. The physical length was found to correlate strongly with the estimated length ($r = 0.95$ in both cases), although the latter experiment showed a greater compression of the length estimates (slope of the linear regression function equal to 0.78 in the former experiment, and to 0.44 in the latter experiment). An analysis of the relationships existing between the acoustical and perceptual levels was carried on using three acoustical features: signal duration, amplitude, and spectral centroid. Apart from the logarithmic amplitude in the latter experiment, none of the considered acoustical variables predicted the length estimates better than the actual length. Length estimates were finally explained by means of a kinematic analysis of the falling rods. The results of this analysis show potential analogies between the auditory and the tactile domain.

V. Roussarie, S. McAdams, and A. Chaigne. ***Perceptual analysis of vibrating bars synthesized with a physical model*** **[209]**

Roussarie et al. used the MDS methodology to study a set of stimuli, which were synthesized by a physical model of a vibrating bar. The bar density and the damping factor were used as synthesis parameters. All stimuli were equalized in loudness and fundamental frequency. Subjects had to rate, according to an analogical scale, the perceived similarity between paired stimuli. They received no information concerning the nature of the stimuli. Similarity ratings were analyzed with MDS algorithms. A two-dimensional solution was found: The first dimension was well correlated with a power function of the damping factor (a parameter of the physical model) and with a linear combination of the logarithm of the decay time of the amplitude envelope and of the spectral centroid. The second dimension correlated with the bar densities and with the frequency of the second component, that was taken as an estimate of perceived pitch.

1.10 2000

P. A. Cabe and J. B. Pittenger. ***Human sensitivity to acoustic information from vessel filling*** **[46]**

Cabe and Pittenger studied vessel filling by listeners' judgments during different tasks. The first experiment assessed the listeners' ability to distinguish filling events from similar events. Stimuli were generated by pouring water into an open tube. The apparatus was designed so to increase (filling), decrease (emptying), or leaving constant the water level inside the tube during pouring. Subjects were asked to categorize stimuli using these three event-categories. Identification accuracy values ranged from 65% to 87%, depending on the type of event. In experiment 2, subjects were asked to fill the vessel up to the brim or to the drinking level. In the former condition only auditory information was available, whereas in the latter one subjects could use all the available perceptual information (visual, tactile, auditory etc.). Results showed a better performance in the latter condition. Nonetheless, in the auditory condition filling levels were close to the maximum possible level.

In the third experiment, blind and blindfolded subjects were asked to fill to brim vessels of different sizes, and with different water flow velocities. Overall performance was accurate. The vessel size, the flow velocity and their mutual interaction influenced the error (computed as the height of the unfilled portion of the vessel), in a way that it was maximum for the smallest vessel filled with the fastest flow velocity, and minimum for the largest vessel. Conversely, when the error was computed as a percentage of the unfilled vessel height with respect to the vessel total height, then the opposite profile was found. Furthermore, no significant differences between blind and blindfolded participants were found.

Patterns of change specifying the time remaining before the end of an event have been defined as τ [149]. A similar variable may be used by listeners to judge the time remaining for a vessel to be filled to the brim. The final experiment tested this hypothesis. Blindfolded participants were asked to judge the time required for a vessel to be filled to the brim using the sounds generated by filling a vessel to three different levels, with three different flow rates. Responses consistently varied with the filling level, and the estimated filling time was strongly correlated with the actual filling time, this revealing an effective use of the τ variable during the execution of the task.

R. L. Klatzky, D. K. Pai, and E. P. Krotkov. *Perception of material from contact sounds* [137]

Klatzky et al. investigated material discrimination in stimuli with variable frequency and decay modulus τ_d. In the first two experiments subjects had to judge on a continuous scale the perceived difference in the material of an object. Stimuli had the same values of frequency and τ_d, but in the second experiment they were equalized by overall energy. As results did not differ significantly in the two experiments, it could be concluded that intensity is not relevant in the judgment of material difference. Experiments 3 and 4 were conducted on the same set of stimuli used in experiment 2. In the former subjects had to judge the difference in the perceived length of the objects, in the latter they had to categorize the material of the objects using four response alternatives: rubber, wood, glass and steel. Results indicated that judgments of material difference and of length difference were significantly influenced by both τ_d and frequency, even though the contribution of the decay parameter to length difference was smaller than that to material difference. An effect of both these variables was found in a categorization task: for lower decay factors steel and glass were chosen over rubber and plexiglass. Glass and wood were chosen for higher frequencies than steel and plexiglass.

A. J. Kunkler-Peck and M. T. Turvey. *Hearing shape* [142]

Kunkler-Peck and Turvey investigated shape recognition from impact sounds. In the first experiment stimuli were generated by striking steel plates of constant area and variable height/width with a steel pendulum. Participants had to reproduce the height and width of the plates by adjusting the position of several bars within a response apparatus. Although dimensions were underestimated, the subjects' performance revealed scaling of a definite impression of the height and width of plates (i.e., definite scaling). Simple regression models were computed using the plates' dimensions or modal frequencies as predictors. Both predictors were highly efficient in motivating the subjects' performance, as regression models were associated to r^2 coefficients which were higher or equal to 0.95. In the second experiment subjects were asked to scale the dimensions of constant-area and variable-height/width plates made of steel, plexiglass and wood. The type of material was discovered to simply modulate the subjects' estimates (i.e., it was associated to a simple additive effect without altering the estimated H/W ratio). Again, the scaling performance was well justified by the theoretical modal frequencies of the plates.

The remaining two experiments were designed to address shape recognition directly. In the third experiment stimuli were generated by striking a triangular, a circular and a rectangular steel plate (the area was kept constant). Shape was correctly classified at a level well above chance. In the last experiment stimuli were generated by striking circular, rectangular and triangular plates made of steel, wood and plexiglass. Subjects were asked to categorize the material as well as the shape. The material was almost perfectly classified, and shape was correctly classified at a level well above chance. A curious tendency of subjects to associate specific geometrical shapes to specific material types was reported (wood with circle, steel with triangle, plexiglass with rectangle).

S. Lakatos. ***A common perceptual space for harmonic and percussive timbres*** **[144]**

This research provides a direct empirical link between timbre perception and sound source recognition. Stimuli produced using sounds from musical instruments, either producing harmonic tones and percussive sounds, were investigated. All stimuli were equalized in pitch and loudness. Eighteen musicians and sixteen non-musicians were asked to rate analogically the timbral similarity of paired stimuli. In three separate sessions subjects had to rate the harmonic set, the percussive set, and a mixed set including harmonic as well as percussive sounds. Responses were analyzed using MDS techniques, as well as clustering procedures. In all cases a musical training did not appear to determine strong differences in the response profiles. For the harmonic set, the MDS analysis revealed a clustering based on the mode of excitation (impulsive vs. continuous). In fact, the first dimension of the MDS space correlated strongly with the logarithmic rise time. Consistently with that, the principal division among stimuli, as computed by the clustering procedure, divided impulsive from continuous tones. Minor divisions were based on the nature of the proximal stimulus, mainly depending on the spectral centroid proximity rather than on features of the source properties. The second dimension of the MDS solution was highly correlated with the spectral centroid.

A three-dimensional MDS solution was derived for the percussive tones. The first dimension was correlated with log rise time, the second with spectral centroid. The third dimension was associated with "timbral richness". Interestingly, the cluster analysis revealed that the stimuli grouped depending on the features of the physical source. For example membranophones clustered together, as well as instruments with wood cavities or instruments made with metallic plates. A two-dimensional MDS solution was derived for the combined set. Dimensions correlated, respectively, with log rise time and with spectral centroid. String instruments were quite overlapped with bar and tubular percussion instruments, probably because of similarities existing in their physical structure. The cluster solution again revealed a grouping based on the similarities in the physical structure of the instruments: Bars, strings and struck tubes were clustered together, as well as wind instruments, drums, and metal plate instruments.

R. A. Lutfi. ***Auditory detection of hollowness*** **[162]**

Lutfi investigated the recognition of hollowness using stimuli synthesized according to the equation describing the motion of a clamped bar. The equation parameters of density and elasticity were chosen in order to model iron, wood and aluminum bars. During each trial subjects were presented the sound of a hollow and of a solid bar. They were asked to tell which one of the two stimuli had been generated by striking a hollow bar. The inner radius of the hollow bar was chosen in order to keep a performance level between 70% and 90% correct. The bar length was randomly chosen from a normal distribution with a specific mean (10 cm for iron and aluminum, 25 cm for wood bars) and a standard deviation of 0.5 cm. Feedback on the correctness of the response was given after each trial. An analysis of the decision weights revealed two strategies used by different listeners to perform the task. One group adopted a decision strategy based on partial frequencies and decay times, that allowed optimal discrimination between hollow and solid bars. The other group adopted a decision strategy based only on frequency. The effect on performance of a limited processing resolution of the acoustical cues was then analyzed. In this way it was shown that an optimal decision strategy yielded, at best, a small advantage over the decision rule based on frequency.