Journal of Experimental Psychology:
Human Perception and Performance
2010, Vol. 36, No. 2, 462–476

0096-1523/10/$12.00 DOI: 10.1037/a0018388

Integration of Acoustical Information in the Perception of Impacted Sound Sources: The Role of Information Accuracy and Exploitability

Bruno L. Giordano
McGill University

Davide Rocchesso
IUAV University of Venice

Stephen McAdams
McGill University

Sound sources are perceived by integrating information from multiple acoustical features. The factors influencing the integration of information are largely unknown. We measured how the perceptual weighting of different features varies with the accuracy of information and with a listener's ability to exploit it. Participants judged the hardness of two objects whose interaction generates an impact sound: a hammer and a sounding object. In a first discrimination experiment, trained listeners focused on the most accurate information, although with greater difficulty when perceiving the hammer. We inferred a limited exploitability for the most accurate hammer-hardness information. In a second rating experiment, listeners focused on the most accurate information only when estimating sounding-object hardness. In a third rating experiment, we synthesized sounds by independently manipulating source properties that covaried in Experiments 1 and 2: sounding-object hardness and impact properties. Sounding-object hardness perception relied on the most accurate acoustical information, whereas impact-properties influenced more strongly hammer hardness perception. Overall, perceptual weight increased with the accuracy of acoustical information, although information that was not easily exploited was perceptually secondary, even if accurate.

Keywords: auditory cognition, sound source perception, information integration, perceptual weight, impact sounds

Supplemental materials: http://dx.doi.org/10.1037/a0018388.supp

The survival of an animal depends on its ability to exploit sensory information to perceive events and objects in the environment. Common sense suggests that this is possible in the visual, but not in the auditory, domain: "we hear sounds, but we see things" (Rosenblum, 2004, p. 221). Within this view, audition is considered an alerting system that signals the occurrence and approximate location of out-of-sight events, eventually triggering further visual and manual exploration (cf. Van Valkenburg & Kubovy, 2004). A growing literature challenges this view, documenting the ability of untrained listeners to recognize a variety of nonspatial properties of the sound source (e.g., whether a glass bottle bounces or shatters on the floor, Warren and Verbrugge, 1984). These abilities demonstrate a perceptual knowledge of the acoustical specification of the sound source. An issue fundamental to the development of a theory of sound source perception thus concerns the nature of such knowledge, i.e., the nature of the relevant acoustical information and the principles that govern the mapping from acoustical information to perceptual response.

The processes involved in the perception of the sound source have been described in terms of a loop (Pastore, Flint, Gaston, & Solomon, 2008): the mechanics of the sound source structures the acoustical signal (e.g., Fletcher & Rossing, 1991) and the sound

Bruno L. Giordano and Stephen McAdams, Centre for Interdisciplinary Research in Music Media & Technology (CIRMMT), Schulich School of Music, McGill University; Davide Rocchesso, Department of Art and Industrial Design, IUAV University of Venice.

Portions of this research were described in a dissertation submitted to the University of Padova, Italy, in partial fulfillment of Bruno L. Giordano's PhD degree, and some were reported at the Stockholm Music Acoustics Conference in Stockholm, Sweden, in August 2003, at the Auditory Perception, Cognition, and Action Meeting in Vancouver, British Columbia, Canada, in November 2003, and at the Meeting of the Acoustical Society of America in Salt Lake City, Utah, in June 2007. Bruno L. Giordano was partially supported by the European Commission's Future & Emergent Technologies collaborative R&D programme Project number IST-2000 25287 (SOb–The Sounding Object), a Marie Curie Training Site Fellowship hosted by Music Acoustics, Department of Speech, Music and Hearing (TMH), Royal Institute of Technology (KTH), Stockholm, Sweden, FP5 contract number HPMT-CT-2000-00119, and Stephen McAdams' Canada Research Chair in Music Perception and Cognition. Portions of this work were supported by grant RGPIN 312774-05 from the Natural Sciences and Engineering Research Council of Canada to Stephen McAdams. We thank Roberto Bresin, Alain de Cheveigné, and Federico Avanzini for their help with the study of the database of sound sources, with the acoustical analyses, and with the sound synthesis model. We thank Stephen Handel, Bruno Repp, Roberta Klatzky, Massimo Grassi, John Neuhoff and Richard E. Pastore for their comments on earlier versions of this manuscript. The robust analyses presented in this article were carried out using the algorithms provided by Verboven and Hubert (2005).

Correspondence concerning this article should be addressed to Bruno L. Giordano, Schulich School of Music, McGill University, 555 Sherbrooke Street West, H3A 1E3, Montréal, Québec, Canada. E-mail: bruno.giordano@music.mcgill.ca

properties are used to perceive the mechanical properties of the sound source (e.g., Carello, Wagman, & Turvey, 2003; McAdams, 1993). The perceptual process requires at least two decisions: first about which acoustical properties are to be taken into account (e.g., sound frequency), and second about how acoustical information should be weighted perceptually (e.g., the log perceived height of a struck plate approximately doubles with an eight-fold increase in log sound frequency, see Kunkler-Peck and Turvey, 2000, for details). Both of these decisions constitute a perceptual criterion, frequently acquired as a result of one's own interactions with the environment (e.g., Gibson & Pick, 2000).

Two opposing views of the development of perceptual criteria emerge from the literature on perception at large. They disagree on whether sensory information that accurately specifies the state of matters in the environment is primary for earlier rather than later stages of development. Followers of the ecological approach posit a progressive increase of the perceptual relevance of the most accurate information (e.g., Gibson & Pick, 2000). Perceptual development would thus proceed toward an increase in the specificity of information, and the adult perceiver would focus on the acoustical feature that uniquely specifies a property of the sound source, i.e., an acoustical invariant (e.g., Michaels & Carello, 1981). The acoustical invariant for the size of a struck object would be, for example, in a one-to-one correspondence with the size itself and would not be influenced by variations in other properties of the sound source (e.g., material). In contrast, Kellman (1996) hypothesizes that development leads to a decrease in the specificity of the information for perception. As such, young perceivers emphasize the avoidance of perceptual errors and focus on the most accurate information. Adult perceivers are instead more concerned with the robustness of a perceptual process that should operate even in the absence of the most accurate information. Consequently, at later stages the perceiver exploits both accurate and less accurate information (Kellman, 1996). Interestingly, a theory of invariance alludes to a dichotomy between perfectly accurate and perfectly inaccurate information (e.g., an acoustical invariant for size vs. a sound property unaffected by size). On the other hand, the position of Kellman (1996) implies a continuum of information accuracy, which also contemplates intermediate cases of less-than-perfectly accurate information (e.g., an acoustical feature that increases with object size for a large number of different objects, but not for all of them).

Consistent with the position of Kellman (1996), empirical investigations on source perception have demonstrated that adult listeners integrate information from multiple acoustical features. The geometry of impacted objects is discriminated perceptually using the frequency of the spectral components of a sound (Houix, 2003; Lakatos, McAdams, & Caussé, 1997; Lutfi, 2001) and the structure of the sound decay (Houix, 2003; Lutfi, 2001) or the spectral center of gravity (Lakatos et al., 1997), the main acoustical determinant of brightness (Grey & Gordon, 1978; McAdams, Winsberg, Donnadieu, De Soete, & Krimphoff, 1995). Perception of the material of impacted objects is influenced by the decay of a sound (Avanzini & Rocchesso, 2001; Giordano & McAdams, 2006; Klatzky, Pai, & Krotkov, 2000; Lutfi & Oh, 1997; Roussarie, 1999) and by its frequency (Avanzini & Rocchesso, 2001; Giordano & McAdams, 2006; Klatzky et al., 2000). Estimation of the hardness and size of a striking object relies on loudness- and brightness-related information (Freed, 1990; Grassi, 2005). The gender of a clapper and of a walker is identified using rate and spectral shape information (Li, Logan, & Pastore, 1991; Repp, 1987). Perception of the length of a rod bouncing on the floor is influenced by the amplitude, frequency and energy decay of a sound (Carello, Anderson, & Kunkler-Peck, 1998).

Despite our knowledge of the acoustical determinants of source perception, the principles underlying the integration of acoustical information are largely unknown. The concept of perceptual weight is central to this matter. When a listener estimates the sound of a struck object focusing on both sound frequency and duration, a perceptual weight is assigned to each of these sound properties. For example, if pitch is weighted more heavily than duration, two sounds are perceived as generated by objects of vastly differing sizes even when their difference in pitch is small, whereas a very large difference in duration will be required to produce the same difference in perceived size. Thus, the task of understanding the integration of information becomes that of unraveling the principles that govern the assignment of perceptual weights to sound properties. Two factors have a potential influence on this process: firstly, the accuracy of the acoustical information within the environment in which the perceptual criteria develop; secondly, the ability of a perceptual system to exploit the acoustical information.

The construct of information accuracy measures the extent to which levels of a source property are reliably diversified by levels of a sound property within the learning environment. Source perception criteria are likely acquired by repeatedly comparing sound properties (e.g., is the sound decaying slowly?) with the estimates of a source property based on nonauditory information: visual (is the object transparent?), tactile (is the object hard and cold?), verbal ("this is glass"), and context-related (e.g., "I am in a restaurant and glasses are likely to clink"; Ballas and Mullins, 1991). Independent of the details of the learning process, across the repeated experiences of a sound source property (e.g., material), acoustical features will differ in the extent to which they reliably diversify its levels (e.g., all the experienced metal sounds have a longer duration than all the experienced plastic sounds, while both materials produce sounds of similar loudness; cf. Giordano & McAdams, 2006). Information accuracy can thus be defined in probabilistic terms, taking into account the particular source property and the particular perceptual task. For example, the accuracy of duration-related information for material identification can be given by the percent overlap between the distributions of durations for sounds of different materials, across our previous experiences of this property. Also, if the task is to rate the hardness of an object, information accuracy can be given by the absolute value of the correlation between values of the physical hardness and values of a specific acoustical feature. Based on previous hypotheses on the perceptual weight of accurate information (Gibson & Pick, 2000; Kellman, 1996), we then expect that a listener will weight acoustical information in proportion to its accuracy. For example, if frequency specifies the size of an object more accurately than sound level, perceptual estimation of size will weight frequency more heavily than level.

Another factor potentially influencing the structure of perceptual criteria is the ability to exploit the information carried by different acoustical features. At least three factors potentially affect information exploitability: discrimination, memory and learning abilities. The ability to discriminate different types of information is considered to be a good predictor of their perceptual weight in both multisensory and unimodal contexts (Ernst & Banks, 2002; Ernst & Bülthoff, 2004). In particular, a perceiver assigns heavier perceptual weights to sensory information that is more easily discriminated (see Lutfi and Liu, 2007, for the effects of sensory noise on perceptual performance). Also learning and memory likely constrain our ability to exploit sensory information. Notably, superior discrimination does not imply better learning and memory abilities. For example, while long-term retention of pitch is surprisingly accurate only in absolute pitch (AP) possessors (Bachem, 1954), AP listeners do not have superior auditory acuity compared to non-AP musician listeners (e.g., Fujisaki & Kashino, 2002; Levitin & Rogers, 2005). Here, we focus on the general ability of a listener to exploit acoustical information, without investigating the differential role of discrimination, memory and learning abilities. In line with previous studies on multisensory integration (e.g., Ernst & Banks, 2002), we expect that independently of the task at hand, a listener will weight more heavily acoustical information that is more easily exploited.

In this study, we investigate the extent to which the perceptual weighting of acoustical information is modulated by its accuracy and exploitability. We focus on a sound source that densely populates the everyday acoustical environment: the impacted sound source (Ballas, 1993). An impact sound is generated by a brief interaction between two objects (e.g., a wooden spoon tapping on a glass), which generates an impulsive acoustical signal with a sharp onset and a more or less abrupt decay. The study of the perception of impacted sound sources has invariably adopted the hammer-sounding object paradigm, according to which the hammer, a highly damped object, impacts the sounding object which is relatively free to vibrate (e.g., Lakatos et al., 1997). Given its higher damping, the shorter and weaker signal radiating from the hammer is likely to be masked by the signal radiating from the sounding object (e.g., Moore, 2003, pp. 65–126). Nonetheless, the acoustical properties of the impact sound are influenced by the properties of the hammer, along with those of the sounding object and those of the impact itself, such as the duration of the contact between the two objects during the impact (Chaigne & Doutaut, 1997; Fletcher & Rossing, 1991). For example, heavier hammers generate louder sounds, whereas smaller and stiffer sounding objects generate sounds with a higher pitch and a longer duration.

The properties of the impact (e.g., contact time) are a particularly appropriate ground for ascertaining the perceptual role of information accuracy. The impact properties and the sound properties they influence are indeed co-determined by the mechanics of the hammer and the sounding object (e.g., contact time and auditory brightness decrease with an increase in the stiffness of either the hammer or the sounding object, Chaigne & Doutaut, 1997). As such, a listener who focuses on impact properties when instead asked to estimate the properties of the hammer or of the sounding object will likely be basing his or her judgment on inaccurate acoustical information. For several reasons, the evidence concerning the perceptual relevance of impact properties is not very compelling. First, several studies carried out on simulated sound sources did not manipulate the impacts (Avanzini & Rocchesso, 2001; Klatzky et al., 2000; Lutfi & Oh, 1997). Second, studies conducted on real sound sources manipulated either the material of the hammer (Freed, 1990) or the sounding object (Giordano & McAdams, 2006; Kunkler-Peck & Turvey, 2000; Tucker & Brown, 2003), but not both. As a result, impact and material properties were in a perfect monotonic relationship and thus statistically equivalent. Third, even when they were manipulated orthogonally to those of the sounding object (Roussarie, 1999), they were varied across experimental sessions, thus decreasing the likelihood of a perceptual effect. Finally, when all these conditions were not met, statistical tests to ascertain their relevance were simply not carried out (Giordano, 2003; McAdams, Kudo, & Kirchner, 1998).

With this study, we investigate the extent to which the accuracy and exploitability of acoustical information influence the perceptual estimation of hammer and sounding-object hardness. We compute measures of information accuracy from the study of a large database of real impacted sound sources. In Experiment 1, two independent groups of participants were trained in the discrimination of hammer or sounding-object hardness. Measures of the exploitability of the most accurate information are derived from measures of discrimination performance. In Experiment 2, the same trained listeners and an additional group of untrained listeners rated the hardness of hammer or sounding object on a different set of real impact sounds. They received no feedback on accuracy. We measured the extent to which hardness estimation performance was significantly influenced by training. The extent to which previous training facilitates performance gives an additional measure of the ability of a trained listener to exploit accurate acoustical information. In Experiment 3, we finally provide more detailed evidence for the effects of information accuracy and hammer/sounding object impact on hardness estimation. To this purpose, we employ a sound synthesis model that allows us to decorrelate acoustical features that covaried in Experiments 1 and 2.

Acoustical Information in the Impacted Sound Source

We collected measures of the accuracy of acoustical information from the analysis of a large database of impacted sound sources, comprising variations in the hardness of the hammer and sounding object, in the sounding object size, and in the impact properties (for details, see Giordano, 2005, pp. 132–148). The database was taken as approximating the variety of impacted sound sources encountered under everyday conditions, and thus the learning environment of untrained listeners. We extracted several acoustical features from each of the sounds. Information accuracy was quantified with reference to a hammer or sounding-object hardness rating task, used in Experiments 2 and 3, by simulating the performance that an ideal listener could reach by focusing, in turn, on each of the sound features. For the sake of completeness, we also measured the extent to which the same acoustical features constituted accurate information for the simulated ratings of impact properties and sounding object size.

Method

We quantified several mechanical properties of the impacted sound sources, and extracted an ordinal measure of the hardness of the different materials. This hardness measure was used to quantify the hardness estimation performance of participants in Experiments 1–2. The sounds in the database were characterized with a set of 10 acoustical descriptors. These descriptors were used to derive measures of information accuracy and to quantify the acoustical criteria used by participants in the behavioral experiments.

Sound Sources

All sound sources were composed of the sounding object, a freely vibrating square plate (thickness = 1 cm), and a hemispherical hammer (radius = 1 cm). The plates and hammers were made of seven different materials: aluminum, alumina ceramic, soda-lime glass, oak, pine, polymethyl methacrylate (plexiglas), and steel. For each material, plates had three areas: 225, 450, and 900 cm^2. For logistical reasons, the two smaller steel plates were made of a less stiff material ($steel_1$) than the largest plate ($steel_2$). Hammers were manufactured using the same seven materials as for the plates ($steel_2$ excluded).

We measured several mechanical properties of the impacted sound sources (for details, see Giordano, 2005, pp. 132–148, and Supplemental Online Materials). Each of these properties belonged to one of four categories: (1) material-related properties of the sounding object (density and elastic coefficients, where metals are denser and stiffer than woods, i.e., they are characterized by higher values of the elastic coefficients); (2) material-related properties of the hammers (density and elastic coefficients); (3) the size of the sounding object; (4) properties of the hammer/sounding-object impact (the force stiffness coefficient *K*, a measure of the extent to which the hammer is compressed during the impact as a result of a given striking force, the maximum impact force and hammer acceleration, and the hammer/sounding-object contact time; a metal hammer produces a greater acceleration, striking force and *K*, and a shorter contact time than would a wood hammer).

We computed the robust (i.e., outlier-independent) Spearman rank correlation ρ between mechanical descriptors within each of the four above-mentioned categories. We used the Minimum Covariance Determinant (MCD) method (Rousseeuw & van Driessen, 1999) to detect outliers. Strong positive correlations were present between the material-related properties of the sounding object or between the material-related properties of the hammer, grand-average $\rho = .92$; $SD = .07$; $df \geq 124$, and also between the impact properties, average $|\rho| = .98$; $SD = .01$; $df \geq 125$. In particular, stiffer materials were also denser (cf. Waterman & Ashby, 1997), and longer hammer/sounding-object contact times were observed for lower hammer accelerations, impact forces and values of the stiffness coefficient *K*.

We separately reduced the hammer and sounding-object material-related mechanical descriptors and the impact variables into three different ordinal variables. To this purpose, we carried out one robust Principal Components Analysis (PCA; Hubert, Rousseeuw, & Branden, 2005) on the rank-transformed mechanical descriptors of each of the three groups. For each group, the final reducing variable was the first Principal Component (PC), which was always strongly correlated with the original rank variables, grand-average $|\rho| = .98$; $SD = .02$; $df \geq 124$. Note that the correlation between the reduced variables and the respective PC was always positive (e.g., density increased for increasing values of the sounding-object material PC), except for the contact time, which decreased for increasing values of the impact PC. The hammer and sounding-object material-related PCs were taken as an ordinal measure of the hardness of the materials. Materials were ordered as follows from softer to harder: pine, plexiglas, oak, ceramic, glass, aluminum, steel1, steel2. Finally, both the hammer and sounding-object hardness PCs were positively correlated with the impact-properties PC, $\rho(135) = .57$; $p < .01$, and $\rho(134) = .79$; $p < .01$, respectively. Shorter contact times were observed with harder hammers and sounding objects (cf. Chaigne & Doutaut, 1997).

Recording Session

We recorded 147 impact sounds, striking each of the 21 plates (seven materials and three areas) with each of the 7 hammers. Sounds were generated in an acoustically isolated room. Plates were mounted in a wooden device by means of nylon threads; hammers were mounted on an aluminum guide heavily damped with a heavy piece of cloth (see Figure 1). Plates were stabilized by means of side weights, so as to avoid drastic movements after the hammer stroke. Plates were struck in their centers by releasing the guide from a fixed angle. A Brüel & Kjær type 4003 condenser microphone was positioned 25 cm from the center of the plate opposite the struck surface. The signal captured by the microphone was delivered to a Symetrix SX202 microphone preamplifier, connected to a Loughborough Sound Images PC/C32 DSP board. The signal was acquired through the DSP board with a sampling rate of 44100 Hz and a resolution of 16 bits.

Acoustical Features

Sounds were analyzed in terms of continuous acoustical features. Defining an acoustical feature corresponds to formulating a

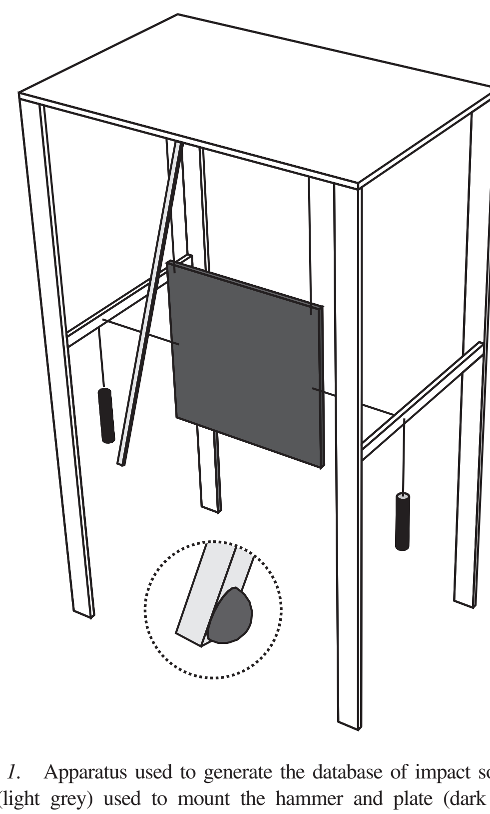

Figure 1. Apparatus used to generate the database of impact sounds. The guide (light grey) used to mount the hammer and plate (dark gray) and the weights (black) used to stabilize the plate after the impact are shown. The bottom part of the figure details how the hammer was mounted on the guide.

hypothesis concerning which properties of a sound signal are perceptually relevant. It should be noted that the problem of acoustical features can be virtually endless, because a large number of mathematical, statistical and signal processing operations can be combined to represent a sound with a single number. Furthermore, even if a first acoustical feature explains the behavioral data perfectly, it is always possible that the perceptual process actually operates on a second, unmeasured feature that is strongly correlated with the first. Addressing these fundamental psychophysical problems is beyond the scope of this study. Here, acoustical descriptors are defined by taking into account the basic properties of the human auditory system, acknowledging the time-varying nature of impact sounds, and focusing on features easily interpreted in terms of basic auditory attributes (loudness, duration, brightness, and pitch). Given the focus of the current study on the perception of the hardness of impacted objects, we also considered acoustical variables strongly associated with the auditory perception of materials in previously published studies (e.g., Giordano & McAdams, 2006; Klatzky et al., 2000).

We extracted 10 acoustical features for each of the sounds by using the methodology detailed in Giordano and McAdams (2006) and in the Supplemental Online Materials. The frequency of the lowest spectral component, F, was estimated from the Fast Fourier Transform of the signal. The remaining nine features were extracted from a simulation of the signal processing stages that take place in the peripheral auditory system. From this representation, we computed the time-varying loudness and spectral center of gravity (SCG), this latter capturing the auditory attribute of brightness (Grey & Gordon, 1978). In general, with impact sounds loudness and SCG decrease monotonically from the peak onset value (see Figure 1 in Supplemental Online Materials). Sound duration, Dur, was defined based on a fixed threshold value for loudness. For both loudness and SCG, we extracted the attack value (initial 10 ms), Lou_{att} and SCG_{att}, and the average value across the signal duration, Lou_{mea} and SCG_{mea}. Three additional descriptors described the decay rate of loudness and SCG. They were operationalized as the slope of the least squares line fitted to specific sections of the temporal signal (rapidly decaying signals yield lower slope values). Lou_{sl1} and Lou_{sl2} measured the rate of loudness decay in the initial and final portion of the sound, respectively. SCG_{slo} measured the rate of SCG decay in the initial portion of the sound. A final descriptor, $\tan\varphi_{aud}$, measured approximately the rate of energy decay in the most intense spectral components of a sound. This descriptor is significantly associated with the auditory identification of impacted solids: more damped materials such as wood and plastics are identified in rapidly decaying sounds where $\tan\varphi_{aud}$ has a higher value (e.g., Giordano & McAdams, 2006).

Results and Discussion

Information accuracy was estimated with reference to a rating task and was given by the absolute value of the robust Spearman rank correlation between acoustical features on the one hand and source properties reduced to independent PCs on the other. Robust bivariate correlations were calculated considering only the *N* non-outliers (minimum $N = 130$, average $N = 140$, $N\ SD = 5.3$).[1] The results of this analysis are presented in Table 1. Note that an information-accuracy score of one indicates that an ideal listener could perfectly order the levels of a source property based on a given acoustical parameter. An information-accuracy score of zero indicates instead that the simulated ratings based on a specific acoustical feature are independent of the actual values of the sound source property. As such, we interpret an information-accuracy score of zero as corresponding to the chance-level performance of an ideal listener who produces random estimates, whereas an information-accuracy score significantly higher than zero is interpreted as indicating better-than-chance rating performance.

We finally isolated the most accurate acoustical feature for the rating of each source property. To this purpose, we tested for each source property whether each of the acoustical features had an information-accuracy score significantly lower than the highest observed score. Significance tests were carried out within a bootstrap framework (Efron & Tibishirani, 1993), based on bias-corrected and accelerated BC_a confidence intervals, a more accurate alternative to the bootstrap percentile method (Efron & Tibishirani, 1993). For each test, 10,000 bootstrap samples were drawn by resampling with replacement from the nonoutlying data points for the contrasted bivariate correlations. The critical *p* value for the rejection of the null hypothesis (the information accuracy score equals the maximum observed value for a given source property) was adjusted in isolation for each of the source properties using the Bonferroni criterion.

F had the highest information accuracy for sounding-object size, $p \leq .002$ across all contrasts, $\tan\varphi_{aud}$ for sounding-object hardness, $p \leq .003$, and SCG_{att} for properties of the hammer/sounding-object impact, $p < .001$. For the hardness of the hammer, two acoustical features yielded the same (maximum) score of information accuracy, Lou_{att} and Lou_{sl1}, $p = .28$.

Consistently with the notion of a continuum of information accuracy, accuracy scores varied widely across acoustical features. Thus, for each of the source properties, several acoustical parameters allowed less-than-perfect but better-than-chance estimation of the mechanical parameters. For example, the best acoustical specifier of the impact properties, SCG_{att}, allowed above-chance estimation of both hammer and sounding-object hardness.

Experiment 1

In this experiment, listeners were trained in a hammer or sounding-object hardness discrimination task. The amount of training necessary to reach target performance level and the performance level reached at the end of the training were interpreted as an estimate of the ability of listeners to exploit the most accurate acoustical information for the perception of hammer and sounding-object hardness. We derived quantitative measures of the influence of acoustical features on behavioral responses, i.e., the perceptual weight of sound properties. For this experiment only, information-accuracy scores were derived directly from the experimental stimulus set, i.e., from the learning context, and not from the database. A heavier perceptual weighting was expected for more accurate acoustical information.

[1] Note that throughout the article, *p* values for the robust correlations are computed considering only the nonoutlying data points. As such, the significance tests for two independent robust correlations computed for starting datasets of the same size might be based on a different number of degrees of freedom.

Table 1
Information-Accuracy Scores for the Acoustical Features Considered in This Study

	Source property					
	Database				Experiment 1	
Acoustical feature	$Size_{SO}$	$Hard_{SO}$	$Hard_{H}$	Impact	$Hard_{SO}$	$Hard_{H}$
$\tan\varphi_{aud}$	$-.22^{**}$	$-.94^{**}$	$-.08$	$-.77^{**}$	1.00^{**}	0.52
Dur	$.55^{**}$	$.89^{**}$	.06	$.68^{**}$	1.00^{**}	0.52
F	$-.74^{**}$	$.43^{**}$	.07	$.25^{**}$	0.56	0.56
Lou_{att}	$.53^{**}$	$-.35^{**}$	$.70^{**}$	$.41^{**}$	0.52	0.92^{**}
Lou_{mea}	.04	$-.81^{**}$	$.21^{**}$	$-.35^{**}$	0.76^{**}	0.68^{*}
Lou_{sl1}	$-.10$	$.75^{**}$	$-.66^{**}$	.01	0.80^{**}	0.72^{**}
Lou_{sl2}	$.57^{**}$	$.86^{**}$	.10	$.61^{**}$	1.00^{**}	0.52
SCG_{att}	.11	$.83^{**}$	$.49^{**}$	$.96^{**}$	0.88^{**}	0.60
SCG_{mea}	$-.55^{**}$	$.73^{**}$	$.22^{**}$	$.63^{**}$	0.76^{**}	0.52
SCG_{slo}	$.24^{**}$	$.80^{**}$	.16	$.46^{**}$	0.80^{**}	0.68^{*}

Note. Hard = hardness; SO = sounding object; H = hammer; Dur = duration; F = frequency; Lou = loudness; SCG = Spectral Center of Gravity; att = attack; mea = mean; sl1 = initial slope; sl2 = final slope; slo = slope. For the database of sound sources, information accuracy is given by the absolute value of the robust Spearman correlation between sound features and source properties. For completeness, the sign of the correlations is also reported. For Experiment 1, information accuracy is given by the proportion of correct discrimination responses afforded by each acoustical feature in a block of 25 2I-2AFC trials.
* $p < .05$. ** $p < .01$. For correlations, $df \geq 128$.

Method

Participants. Twenty-four individuals took part in the experiment (17 women, 7 men; age = 18–58; mean age = 32). They were paid for their participation. Normal hearing was assessed, measuring hearing threshold in both ears at octave-spaced frequencies (125–8000 Hz) using a standard audiometric procedure. Hearing thresholds never exceeded normative values by more than 15 dB (ISO 389–8, 2004; Martin & Champlin, 2000).

Stimuli and apparatus. We selected 15 stimuli from combinations of three levels of material hardness each for the hammer and the sounding object (low: oak; intermediate: ceramic; high: steel) and two levels for the sounding-object size (small: roughly 225 cm^2, and large: roughly 900 cm^2; see Supplemental Online Materials). Sound stimuli were reduced to a duration of 1 s, with a linear offset ramp of 5 ms.

Stimuli were stored on the hard disk of a Macintosh G5 Workstation, equipped with an M-Audio Audiophile 192 S/PDIF interface for digital-to-analog conversion. Audio signals were amplified with a Grace Design m904 monitor system and presented binaurally through Sennheiser HD280 headphones. Participants sat inside an IAC double-walled soundproof booth. Signal peak level at the headphone ranged from 55- to 70-dB SPL as measured with a Brüel & Kjær Type 2205 sound level meter coupled with a Brüel & Kjær Type 4153 artificial ear.

Design and procedure. Participants performed a 2I-2AFC hardness discrimination task. On each trial, they were presented two sounds, and they had to decide in which sound the hammer was harder (H condition) or the sounding object was harder (SO condition).

The paired sound sources within each of the trials differed in the size and material of the sounding object and in the material of the hammer, for a total of 25 different pairs. For the design of this experiment, the 225- and 900-cm^2 steel plates were considered as being made of the same material. The set of 25 stimulus pairs was presented in block-randomized order. The within-pair offset-to-onset duration was fixed at 100 ms. The within-pair order of the sounds was randomly chosen on a trial-by-trial basis. Feedback on response correctness was given at the end of each trial. The experiment was terminated if proportion correct was greater than .75 at the end of a block of trials. If a participant did not reach the performance threshold within 1 hour, the experiment was terminated and that person's data were not considered.

Thirteen participants were randomly assigned to the H condition; eleven to the SO condition. Three participants in the H condition and one of the participants in the SO condition did not reach the performance threshold within 1 hour. The data from 10 participants in each condition were thus analyzed.

Results

We carried out three different analyses. First, we examined differences in discrimination performance between the H and SO experimental conditions. Second, we investigated the acoustical features used by participants to discriminate hardness levels in each of the experimental conditions and quantified their perceptual weight. Finally, we compared the perceptual weights with measures of the accuracy of the acoustical information.

On average, the 75%-correct hardness-discrimination threshold was reached after a significantly higher number of blocks in condition H compared to condition SO: 4.8 and 1.2 blocks, respectively; unpaired-samples $t(18) = 3.94$; $p < .01$. All participants in the H and SO conditions reached the performance threshold after eleven and two blocks of trials, respectively. Performance in the last block of trials was significantly worse in the H condition compared to the SO condition: 78 and 88% correct, respectively; unpaired-samples $t(18) = -4.00$; $p < .01$. Figure 2 shows the evolution of performance in the population of participants for the two experimental conditions. In order to highlight the rapid increase in performance for the very first trials in the SO condition, performance levels are shown for each group of 10 trials instead of for each block of 25 trials.

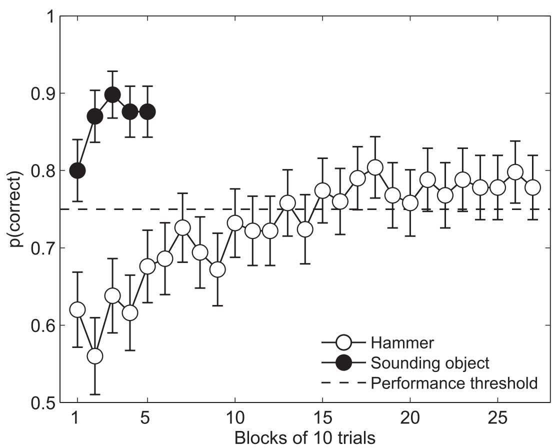

Figure 2. Average hardness discrimination performance in Experiment 1 for hammer and sounding-object judgments. Error bars = ±1 *SEM.*

The analysis of the acoustical correlates of hardness discrimination involved two sequential steps. First, groups of strongly correlated acoustical descriptors were reduced to a single variable. Second, a separate multiple regression model was created for the population data from each of the two experimental conditions using the acoustical descriptors as predictors. The perceptual weights of sound properties were estimated based on these regression models. All steps focused on rank information so as to avoid influences of the shape of the monotone transform relating the different variables on the measure of their association (e.g., Iman & Conover, 1979; Ryan, 1997).[2] We used robust statistical techniques to minimize the influence of outliers on the statistical measures and inferences (Rousseeuw & Driessen, 1999).

Before they were related to behavioral data, the acoustical features were transformed into a set of weakly correlated variables. It is important to note here that the intercorrelations of acoustic parameters related to the different mechanical parameters are often high and it is not possible to decide which is the greatest contributor to the experimental judgments. Their grouping together as a combined descriptor through the approach described below reflects this real-world ambiguity. The data-reduction procedure thus reduced clusters of strongly correlated descriptors by means of PCA. The choice of which descriptors to reduce into the same PC was guided by a cluster analysis carried out on the matrix of correlations among acoustical descriptors across the stimuli.

For the purpose of the current experiment, the acoustical features described the pairs of sounds judged by the listeners and measured the difference between the value of one acoustical descriptor in the first and second sound of a given trial (e.g., the difference between the attack brightness in the first and second sound of the trial). The order of presentation of the sounds within the trial was taken into account for this purpose, giving a total of 50 different pairs of sounds with their corresponding differences for each of the acoustical features. We computed the robust Spearman rank correlation ρ between the acoustical descriptors (Rousseeuw & Driessen, 1999). A hierarchical cluster analysis (average linkage) was carried out on a measure of the distance between acoustical descriptors, defined as one minus the absolute value of their correlation. Finally, a robust PCA was carried out on the ranks of the clusters of strongly correlated acoustical descriptors (Hubert et al., 2005; Tucker, 1960), and the first PC was retained as the final reducing variable. In particular, starting from the condition where each of the descriptors was in an isolated cluster, each of the clusters of descriptors was independently reduced to a single PC, and the correlations between the PCs derived for each of the clusters were computed. If none of the absolute correlations exceeded a threshold of .50, the procedure was terminated. Otherwise, the number of clusters of descriptors was reduced by one and the procedure was iterated.

The data-reduction step yielded two PCs of acoustical descriptors that were weakly correlated with each other, robust $\rho(48) = -.10$; $p = .50$. The two PCs summarized well the acoustical variables they reduced (see Table 2). To facilitate the comparison of results across experiments, the PCs were labeled based on which properties of the sound source were most accurately specified by the sound properties they included. The first of the PCs was labeled Hard_{SO}/Imp because it included both the most accurate specifiers of sounding-object hardness and hammer/sounding-object impact, $\tan\varphi_{\text{aud}}$ and SCG_{att}, respectively. The second PC was labeled $\text{Hard}_{\text{H}}/\text{Size}_{\text{SO}}$ because it included the most accurate specifier of hammer hardness, Lou_{att}, and sounding-object size, F. The complete partitioning of acoustical features into reducing variables is shown in Table 2.

We then measured the extent to which the acoustical PCs predicted population responses in the last block of trials of the hardness discrimination experiment. For each of the experimental conditions, we created one separate multiple robust rank-regression model by rank-transforming both the dependent variable and the predictors, i.e., the acoustical PCs (Iman & Conover, 1979). The dependent variable was the population probability that the first stimulus in a pair was perceived as generated with a harder hammer (H condition) or with a harder sounding object (SO condition). The perceptual weight of the acoustical PCs was estimated using a statistical measure of the size of their effect on the perceptual responses: their partial R^2 (R_p^2) within the multivariate regression model. R_p^2 for a variable X is defined as the ratio of the gain in the proportion of explained variance when X is included in the set of all predictors to the variance left unexplained when X is not in the regression model (e.g., Mulaik, 2005). R_p^2 measures were computed for the observed dependent variable rather than on the rank-transformed values. For both experimental conditions, the number of outliers not included in the final estimation procedure was 3 out of the 50 data points. Both regression models explained more than 83% of the variance in both the observed probabilities and their rank values. Independently of the experimental condition, responses were significantly influenced by the Hard_{SO}/Imp variable, $p < .001$, $df = 44$, whereas $\text{Hard}_{\text{H}}/\text{Size}_{\text{SO}}$ had a significant influence on responses in the H, but not in the SO condition, $p < .001$ and $p = .21$, respectively; $df = 44$. The parameters of the regression models are presented in Table 3.

[2] We used the same nonparametric approach for data from all experiments. Similar conclusions were reached using a parametric approach based on logistic regression models for Experiment 1 (Agresti, 1996) and on linear mixed-effects models (Verbeke & Molenberghs, 2000; West, Welch, & Galecki, 2006) for Experiments 2 and 3.

Table 2
Summary of Data-Reduction Steps for the Quantification of the Acoustical Criteria for Perception

	Experiment 1		Experiment 2		Experiment 3	
Acoustical feature	PC	ρ	PC	ρ	PC	ρ
$\tan\varphi_{aud}$	Har_{SO}/Imp	.96	Har_{SO}/Imp	.97	Har_{SO}	.84
Dur	Har_{SO}/Imp	−.95	Har_{SO}/Imp	−.94	Har_{SO}	−.99
F	Har_{H}/$Size_{SO}$	.94	Har_{H}/$Size_{SO}$	.92	$Size_{SO}$	−.91
Lou_{att}	Har_{H}/$Size_{SO}$	−.96	Har_{H}/$Size_{SO}$	−.94	Imp	−.97
Lou_{mea}	Har_{H}/$Size_{SO}$	−.74	Har_{SO}/Imp	.91	Imp	−.95
Lou_{sl1}	Har_{SO}/Imp	−.78	Har_{H}/$Size_{SO}$	.93	Imp	.91
Lou_{sl2}	Har_{SO}/Imp	−.92	Har_{SO}/Imp	−.95	Har_{SO}	−.97
SCG_{att}	Har_{SO}/Imp	−.93	Har_{SO}/Imp	−.80	Imp	−.85
SCG_{mea}	Har_{H}/$Size_{SO}$	.85	Har_{SO}/Imp	−.86	$Size_{SO}$	−.98
SCG_{slo}	Har_{SO}/Imp	−.90	Har_{SO}/Imp	−.72	Har_{SO}	−.94

Note. PC = principal component; ρ = robust Spearman correlation of acoustical feature with PC; Hard = hardness; SO = sounding object; H = hammer; Imp = hammer/sounding-object impact; Dur = duration; F = frequency; Lou = loudness; SCG = Spectral Center of Gravity; att = attack; mea = mean; sl1 = initial slope; sl2 = final slope; slo = slope. For all correlations, $p < .01$ and $df \geq 38$, 12, and 20 for Experiments 1, 2, and 3, respectively. For each of the acoustical features, and for each of the Experiments, the label of the reducing principal component (PC) is presented along with its robust Spearman correlation ρ with the reducing PC.

Post hoc contrasts were carried out to test for significant within- and between-condition differences in the perceptual weighting of the two acoustical PCs. BC_a bootstrap hypothesis tests were used (10,000 bootstrap samples per contrast, drawn by resampling non-outlying data points). In particular, we contrasted the R_p^2 for different predictors within the same experimental condition (two contrasts) and the R_p^2 for the same predictor in the two experimental conditions (two contrasts). We adopted unidirectional null hypotheses, thus testing whether the difference between the R_p^2 for predictors X and Y was significantly higher than zero if we had observed a higher R_p^2 for X, and significantly lower than zero if we had observed a higher R_p^2 for Y. All contrasts were significant, $p < .001$. Thus, participants focused on $Hard_H$/$Size_{SO}$ in the hammer hardness condition and on $Hard_{SO}$/Imp in the sounding-object hardness condition. Also, the weights given to each of the two PCs differed significantly across conditions.

Finally, we derived information-accuracy scores by measuring the maximum performance that an ideal listener would achieve when performing the 2I-2AFC hardness discrimination task by focusing on a specific acoustical feature for all of the possible 25 pairs of sounds. In particular, the accuracy of the acoustical feature X was defined as the maximum proportion of correct trials that the ideal listener would reach by consistently identifying the paired sound with a higher value of X as harder or softer (e.g., the maximum between the proportion of correct trials reached when always identifying the paired sound with the higher Lou_{att} as harder, and when identifying the paired sound with the higher Lou_{att} as softer). Information-accuracy scores are reported in Table 1. Their significance is calculated using a binomial test. Consistently with the results of the analyses conducted on the database, Lou_{att} and $\tan_{\varphi aud}$ maximized performance in the perception of hammer and sounding-object hardness, respectively. Most importantly, information was present for nearly perfect discrimination of both hammer and sounding-object hardness. Figure 3 plots R_p^2 for the acoustical PCs as a function of the median accuracy scores for the features they reduce. For both experimental conditions, the perceptual relevance of a group of acoustical features increases with its overall accuracy.[3]

Table 3
Summary of Multiple Rank Regression Analysis of Population Responses in the Last Block of Trials in Experiment 1

Variable	*B*	*SE* B	β	R_p^2
		Hammer condition		
Har_{SO}/Imp	−0.40	0.06	$-.41^{**}$	.44
Har_H/$Size_{SO}$	−0.81	0.06	$-.83^{**}$	.80
		Sounding-object condition		
Har_{SO}/Imp	−0.85	0.06	$-.89^{**}$	.93
Har_H/$Size_{SO}$	0.08	0.06	.08	−.11

Note. Hard = hardness; Imp = hammer/sounding-object impact; SO = sounding object; H = hammer; *B* = regression coefficient; β = standardized regression coefficient; R_p^2 = partial R^2. R_p^2 estimates the perceptual weight of the PC of acoustical descriptors.
** $p < .01$; $df = 44$.

Discussion

Consistently with previous source perception studies, hammer hardness discrimination appeared to involve an integration of both the most accurate and least accurate acoustical information for the task. Indeed, the responses of the participants to this condition were consistent with a perceptual focus on features that allowed a high discrimination performance (e.g., Lou_{att}), but also onto features that allowed chance-level performance (e.g., $\tan\varphi_{aud}$). The discrimination of sounding-object hardness was instead influenced by one single source of information, the group of acoustical variables containing the most accurate specifiers of the sounding-

[3] The same trends were observed when considering the measures of information accuracy computed on the database of sound sources, assuming a hardness rating task. Also, the same trends emerged when considering either the maximum or the average of the accuracy scores for the acoustical features reduced by the same PC. This latter consideration applies also to the analysis of data from Experiments 2 and 3.

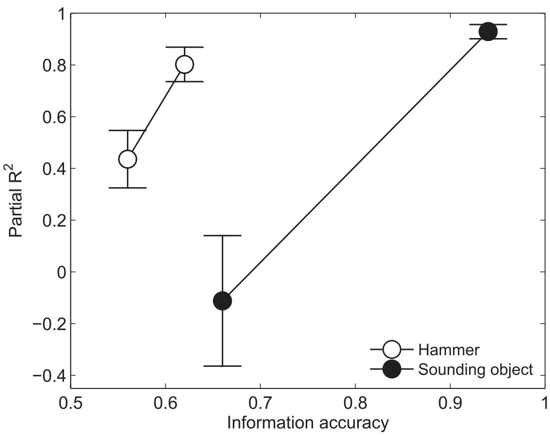

Figure 3. Perceptual weights of the principal components of acoustical descriptors in Experiment 1, as quantified by the R_p^2 measure of effect size, as a function of the median of the accuracy scores for the reduced acoustical features. Error bars = ±1 bootstrap *SEM* (10,000 bootstrap replicates).

object hardness and of the impact properties. This finding is neither consistent nor inconsistent with an integrative perceptual strategy, because it cannot be ascertained whether judgments were focused on one or more of the strongly correlated sound properties.

In line with our expectations, participants in both experimental conditions appeared to weight more heavily the most accurate acoustical information for the task at hand. This finding is particularly relevant, because the same level of discrimination performance can be reached using largely different weighting strategies (cf. Lutfi & Liu, 2007), i.e., better-than-chance perceptual performance does not automatically imply a focus on the most accurate information.

Finally, several measures of discrimination performance suggest an impaired ability to exploit the most accurate acoustical information for the perception of hammer hardness, as compared to the ability to process the most accurate information for sounding-object hardness. Indeed, as compared to the sounding-object condition, participants in the hammer condition reached lower levels of discrimination performance and required longer training to reach the target performance level. Furthermore, whereas all participants in the sounding-object condition reached the target performance level, 17% of those in the hammer condition did not. Notably, the same pattern of results can be explained not only by differences in the exploitability of accurate information, but also by differences in pre-experimental expertise with the task. In other words, higher performance levels in the sounding-object condition could have emerged because participants were already more skilled at discriminating sounding-object hardness as compared to hammer hardness before taking part in the experiment. We addressed this alternative hypothesis in Experiment 2.

Experiment 2

In this experiment, we investigated the perceptual estimation of hammer and sounding-object hardness with naive listeners and with the listeners trained in Experiment 1. Participants rated the hardness of either object with a novel set of impacted sounds. They were given no feedback. The comparison of recognition performance for the hammer and sounding-object hardness in untrained listeners allowed us to assess whether the lower hammer discrimination performance observed in Experiment 1 was caused by an expertise with the task that was already lower before the training took place. Further evidence on the exploitability of accurate information for hardness estimation was gathered by comparing the performance of trained and untrained listeners. A training-related improvement of perceptual performance was thus taken as additional evidence for the ability to retain the perceptual criteria acquired during Experiment 1, which focused on accurate information, and to generalize them to a novel set of sounds even in the absence of feedback on response correctness. The perceptual weight of acoustical information was finally compared with the information-accuracy scores computed for the database of impact sounds.

Method

Participants. Forty normal-hearing individuals took part in the experiment (28 women, 12 men; age = 18–40 years; mean age = 22). Among them, 20 had already participated in Experiment 1. They were paid for their participation. Hearing thresholds were assessed using the same procedure as in Experiment 1.

Stimuli and apparatus. We selected 18 stimuli from the database of impacted sound sources (see Supplemental Online Material). None of the selected stimuli had been investigated in Experiment 1. The experimental set included all the sounding-object and hammer materials, and contained small, medium and large plates. Acoustical signals were reduced to a duration of 1 s with a linear offset ramp of 5 ms.

The same apparatus as for Experiment 1 was used for stimulus presentation and data collection. Signal peak levels ranged from 55- to 67-dB SPL, as measured with the same equipment as for Experiment 1.

Design and procedure. Participants rated the hardness of the hammer or the sounding object by moving a slider along a scale labeled "very soft" and "very hard" at the two extremes. Stimuli could be replayed as many times as needed before giving the rating. Before the experimental phase, participants were required to listen to all the stimuli in random order at least three times. This allowed them to get a sense of the within-set variation of hammer or sounding-object hardness. The 18 stimuli were then presented in blocked-randomized order for each of 10 repetitions for a total of 180 trials. No feedback on response correctness was given.

Separate groups of 10 listeners participated in the four experimental conditions: hammer hardness with or without previous discrimination training (H_{tr} and H_{ntr}, respectively) and sounding-object hardness with or without previous discrimination training (SO_{tr} and SO_{ntr}, respectively). The trained listeners participated in this experiment immediately after Experiment 1, following a break of 10 minutes maximum.

Results

We carried out the same analyses as for Experiment 1 involving a quantification of the hardness estimation performance, the measurement of the perceptual weight of acoustical information, and the comparison of perceptual weights with measures of information accuracy.

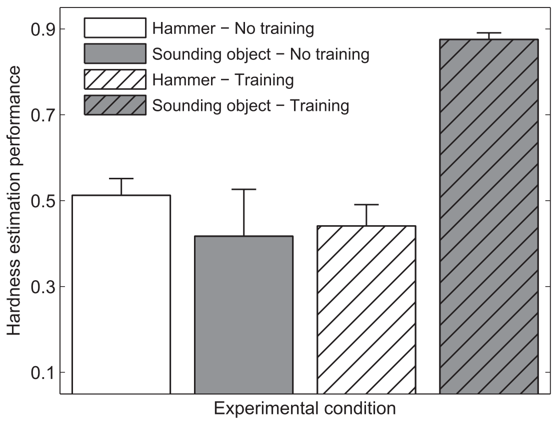

Figure 4. Average hardness estimation performance for the different conditions of Experiment 2. Error bars = ±1 *SEM.*

We quantified the hardness estimation performance for each of the participants on a block-by-block basis. To achieve this, we computed, for each block of trials the robust Spearman correlation between the estimated hardness and the actual hardness derived from the study of the database of impact sounds. Following the approach used to derive information-accuracy scores in the database, correlations of one and zero were interpreted as corresponding to perfect and chance-level performance, respectively. Figure 4 shows the hardness estimation performance for each of the conditions averaged across participants and blocks of trials. Overall, participants in both the H_{ntr} and H_{tr} conditions performed better than chance in 69% of the blocks of trials, whereas participants in the SO_{ntr} and SO_{tr} conditions performed better than chance in 72% and 100% of the blocks of trials, respectively (critical *p* value = .05; $df \geq 12$ for all tests). We further analyzed the performance measures with a repeated-measures ANOVA model, with blocks as within-subject factor and training (trained vs. untrained) and judged object (hammer vs. sounding object) as between-subjects factors. Neither the main effect of blocks, nor any of the interactions including this variable were significant, $F(9, 324) < 1$, $\eta_p^2 \leq .03$; $p \geq .50$ after Huynh-Feldt correction, $\varepsilon = .96$. The interaction between training and judged object was significant, $F(1, 36) = 5.53$; $\eta_p^2 = .13$; $p = .02$, whereas the main effects of both of these factors fell short of significance, $F(1, 36) \leq 2.96$; $\eta_p^2 \leq .14$; $p \geq .06$. Further contrast analyses showed that training improved rating performance in the sounding-object condition, $F(1, 18) = 4.52$; $\eta_p^2 = .20$; $p = .05$, whereas the effect of training in the hammer condition was not significant, $F(1, 18) = 1.20$; $\eta_p^2 = .06$; $p = .29$. Also, estimation performance was significantly higher for the sounding object than for the hammer after training, $F(1, 18) = 49.37$; $\eta_p^2 = .73$; $p < .01$, but not in the absence of training, $F(1, 18) = 1.93$; $\eta_p^2 = .01$; $p = .67$.

The analysis of the acoustical correlates of hardness ratings involved a data reduction step and a subsequent regression modeling step. We adopted the same robust rank-based methodology as for Experiment 1. The data-reduction step was carried out on the acoustical features of the 18 experimental stimuli. Two uncorrelated PCs of acoustical descriptors were extracted, robust $\rho(16) = -.44$; $p = .07$. The reduced acoustical descriptors were strongly correlated with the reducing variables (see Table 2). The first PC was labeled $Hard_H/Size_{SO}$, because it included the most accurate specifiers of hammer hardness, Lou_{att} and Lou_{s11}, and of sounding-object size, F. The second PC was labeled $Hard_{SO}/Imp$, because it included the most accurate specifier of sounding-object hardness, $\tan\varphi_{aud}$, and properties of the hammer/sounding-object impact, SCG_{att}. One robust regression model was computed for each of the four experimental conditions with the acoustical PCs as independent variables. For each of the participants, we initially computed the median of the hardness estimates across the blocks of trials. The across-participants condition-specific median of these estimates was used as the dependent variable. The parameters of the regression models for data from all experimental conditions are presented in Table 4. Across regression models, the maximum

Table 4
Summary of Multiple Regression Analysis of Hardness Estimates in the Population of Participants in Experiment 2

Variable	*B*	*SE* B	β	R_p^2
		Hammer—No training condition		
Har_{SO}/Imp	−0.59	0.09	−.64**	.74
$Har_H/Size_{SO}$	0.34	0.10	.37**	.29
		Sounding-object—No training condition		
Har_{SO}/Imp	−0.71	0.18	−.64**	.66
$Har_H/Size_{SO}$	0.46	0.18	.41*	.56
		Hammer—Training condition		
Har_{SO}/Imp	−0.93	0.15	−.93**	.72
$Har_H/Size_{SO}$	−0.23	0.15	−.23	.11
		Sounding-object—Training condition		
Har_{SO}/Imp	−0.76	0.08	−.80**	.91
$Har_H/Size_{SO}$	0.16	0.08	.17	.00

Note. Hard = hardness; Imp = hammer/sounding-object impact; SO = sounding object; H = hammer; *B* = regression coefficient; β = standardized regression coefficient; R_p^2 = partial R^2. R_p^2 estimates the perceptual weight of the PC of acoustical descriptors.
** $p < .01$; $df \geq 11$.

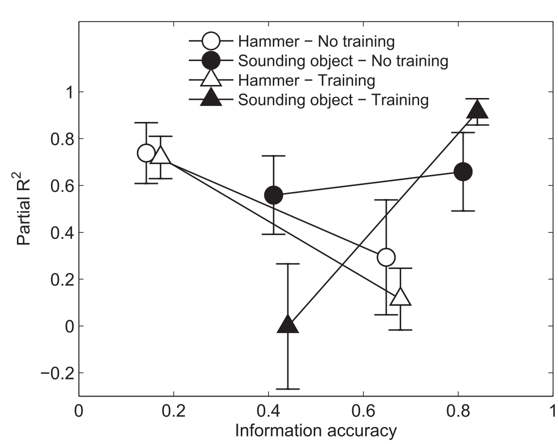

Figure 5. Perceptual weights of the principal components of acoustical descriptors in Experiment 2, as quantified by the R_p^2 measure of effect size, as a function of the median of the accuracy scores for the reduced acoustical features. Error bars = ±1 bootstrap *SEM* (10,000 bootstrap replicates).

number of outliers was 4 out of the 18 data points. On average, the regression models explained 83% of the variance of the hardness ratings and of their ranks ($SD = 10\%$ in both cases).

We tested for significant within- and across-condition differences in the perceptual weight of the acoustical PCs. We adopted the same bootstrap procedure as for Experiment 1, based on the R_p^2 measure of perceptual weight. We computed four within-condition contrasts for pairwise R_p^2 differences between PCs. The between-condition contrasts tested for pairwise differences in the R_p^2 for the same PC, as determined by a change in either the training or in the judged object factors in isolation, but not by a change in both (e.g., we contrasted the R_p^2 for $\text{Hard}_\text{H}/\text{Size}_\text{SO}$ in the H_ntr condition with that in all of the other conditions but SO_tr). In total, we carried out eight between-condition contrasts. We adopted a family-wise Bonferroni correction for all contrasts, $\alpha = .0125$ and .006 for the within- and between-condition contrasts, respectively. Within-condition contrasts revealed that, independently of training, participants in the hammer conditions focused more on $\text{Hard}_\text{SO}/\text{Imp}$ than on $\text{Hard}_\text{H}/\text{Size}_\text{SO}$, $p \leq .008$. The same trend was observed for the sounding-object conditions, although a significantly heavier weighting of $\text{Hard}_\text{SO}/\text{Imp}$ emerged only for participants in the training condition, $p < .001$ and $p = .30$ for the SO_tr and SO_ntr conditions, respectively. None of the between-conditions contrasts was significant, $p \geq .02$. However, weak evidence supported a heavier weighting of $\text{Hard}_\text{SO}/\text{Imp}$ in the SO_tr condition than in the H_tr and SO_ntr conditions, $p = .02$, and a heavier weighting of the $\text{Hard}_\text{H}/\text{Size}_\text{SO}$ in the SO_ntr than in the SO_tr condition, $p = .03$.

Finally, we related the perceptual weight of the acoustical PCs to the information-accuracy scores computed on the database (see Figure 5). For the sounding-object conditions, the perceptual weight of a group of acoustical features increases with its overall accuracy, whereas the opposite trend emerges for the hammer conditions. Finally, it should be noted that an increase of perceptual weight with information accuracy does not appear to be significant for data from the SO_ntr condition, because the pairwise within-condition contrast for the R_p^2 of the two PCs was not significant.

Discussion

In Experiment 1, we observed a comparatively impaired performance for the discrimination of hammer hardness. This result was hypothesized to originate from either a lower exploitability for the accurate acoustical information for the task or from a comparatively lower expertise with hammer perception prior to participation to Experiment 1. This latter hypothesis was disconfirmed in Experiment 2, where untrained listeners performed equally well when estimating the hardness of either the hammer or the sounding object. Consistently with a lower exploitability of accurate hammer hardness acoustical information, we further observed a significant training-related improvement of hardness estimation performance for the sounding object, but not for the hammer. It is likely that the trained participants in the hammer condition were not able to store the perceptual criteria used in Experiment 1 and/or to generalize them to a situation in which feedback was not given and where a different hardness perception task was carried out on a different set of sounds. The analysis of the acoustical criteria for hardness estimation revealed that participants in all conditions focused primarily on the $\text{Hard}_\text{SO}/\text{Imp}$ group of acoustical features. Interestingly, whereas all untrained participants were also influenced by variations in the $\text{Hard}_\text{H}/\text{Size}_\text{SO}$ group of features, this was ignored by all trained participants. Although this result is consistent with a more selective perceptual criterion for trained participants, we should remember that it does not disconfirm the integration of information from multiple acoustical features within the $\text{Hard}_\text{SO}/\text{Imp}$ group. What this result instead proves is that the trained participants who discriminated hammer hardness in Experiment 1 by focusing on $\text{Hard}_\text{H}/\text{Size}_\text{SO}$ ignored this component in Experiment 2.

The comparison of perceptual weights with the measures of information accuracy revealed that, independently of training, the estimation of sounding object hardness focused on the most accurate information, whereas the estimation of hammer hardness did not. This latter result is in sharp contrast with those of Experiment 1, where hammer hardness was also estimated focusing on the most accurate information. Three methodological factors might explain this difference: a change in stimulus set; a change in task (discrimination vs. ratings); the absence of feedback in Experiment 2. The first explanation is unlikely, because the stimuli of both Experiments 1 and 2 comprised similar ranges of variation for each of the properties of the sound source. Further, it is unclear why a change in sound set would affect the criteria for hammer but not for sounding-object hardness perception. This latter consideration applies also for the hypothesis of a change in perceptual criteria caused by a change in task. It is more likely, instead, that participants in Experiment 1 needed trial-by-trial feedback on hammer hardness discrimination performance to sustain a focus on the most accurate, but hardly exploitable information. As such, participants in Experiment 1 were perhaps unable to store in long-term memory the perceptual criteria based on the most accurate information for hammer hardness. As a consequence, when feedback was unavailable in Experiment 2, the same participants reverted to pre-existing perceptual criteria, which likely did not focus on the most accurate information for hammer hardness.

Based on these considerations, the results of Experiment 2 apparently support a joint influence of information accuracy and exploitability on the structure of perceptual criteria. However, a fourth alternative hypothesis can be advanced to explain why the most accurate information for hammer hardness was ignored even by the trained participants in this experiment, i.e., that all participants had a generalized bias towards focusing judgments on the material of sounding-object, the main vibrating object. A rigorous test of this alternative hypothesis was not possible with the real sound sources investigated in Experiment 2 however, because the acoustical specifiers of the sounding-object hardness were strongly correlated with those of the impact properties. We addressed this point in Experiment 3, investigating synthetic sounds generated with a physical model of the impacted sound source.

Experiment 3

We carried out a final experiment on the perceptual estimation of the hammer and sounding-object hardness. Sound stimuli were synthesized using a physically inspired model of the impacted sound source. We could thus independently manipulate the acoustical specifiers of the two properties of the sound source with a potentially strong influence on the perception of the hardness of the two objects: the sounding-object hardness and the properties of the hammer/sounding-object impact. These acoustical specifiers were strongly correlated in both Experiments 1 and 2. By using the synthesis model, we could thus test whether the results of Experiment 2 were produced by a perceptual bias towards focusing on the sounding-object material, in which case we expected a generalized perceptual focus on the sounding-object hardness, or by a joint effect of the accuracy and exploitability of acoustical information, in which case we expected a stronger focus on the properties of the hammer/sounding-object impact for the perception of the hammer hardness.

Method

Participants. Fifty-one individuals took part in the experiment on a voluntary basis (33 women, 18 men; age: 19–53; mean age = 24). All of them reported normal hearing. None of them had participated in Experiments 1 and 2.

Stimuli and apparatus. Stimuli were synthesized using the model of a struck bar (see Supplemental Online Materials). A total of 27 sounds was investigated, given by the factorial combination of three levels for each of three different parameters: the lowest frequency of the signal F, i.e., the best acoustical specifier of the sounding-object size (F = {50, 200, 800} Hz); tanφ, modeling variations in the material of the bar ($\tan\varphi = \{1.99, 7.96, 31.83\} \times 10^{-3}$); the force-stiffness coefficient *K*, modeling variations in the stiffness of the hammer/sounding-object impact ($K = \{0.05, 2.24, 100\}$ $N/m^{1.5}$, corresponding to a duration of the hammer/sounding-object contact during the impact of 2.55, 0.56, and 0.12 ms, respectively).

Stimuli were stored on the hard disk of an Intel-PC workstation equipped with a Sound Blaster Live sound card. Audio signals were amplified with a Sennheiser HEV70 amplifier and presented binaurally through Sennheiser HE60 headphones. Participants sat inside an acoustically isolated room. Signal peak level ranged from 41- to 88-dB SPL as measured with a Brüel & Kjær Type 2238 sound-level meter.

Design and procedure. Participants estimated the hardness of the hammer (H condition) or the sounding object (SO condition) on a scale ranging from 1 (*really soft*) to 100 (*really hard*). They typed a number on the keyboard after the sound was presented. Before giving the response, they were allowed to replay the stimulus as many times as needed. The 27 stimuli were presented in blocked-randomized order for each of 10 repetitions for a total of 270 trials. No feedback on response correctness was given. Data from the first block of trials, meant to familiarize participants with the task and with the within-set variability in perceived hardness, were not analyzed. Prior to the beginning of the experiment, we informally assessed hardness discrimination abilities of the participants. They were asked to identify the hardest hammer (H condition) or the hardest sounding object (SO condition) in two different pairs of sounds generated by striking two different real sounding objects (metal or plastic bowl) with two different hammers (felt or wood). In the H and SO conditions, the sounding-object and hammer materials were kept constant within the pair, respectively. The sounds were generated out of the participant's sight. No feedback on response correctness was given.

We randomly assigned 24 and 27 participants to the H and SO conditions, respectively. Three of the participants in the SO condition failed in at least one of the initial hardness discrimination trials. Their data were not considered.

Results

Following the methodology used in Experiments 1 and 2, we quantified the perceptual weights of acoustical information and compared these measures with the information-accuracy scores extracted from the database.

The data reduction step was carried out on the acoustical descriptors of the 27 experimental stimuli. Three PCs of acoustical descriptors were extracted. Neither the rank correlation between PCs 1 and 3, nor that between PCs 2 and 3 were not significant, $|\rho| \leq .31$; $p \geq .12$; $df \leq 24$. The correlation between PCs 1 and 2 was weak, but significant, $\rho(23) = .44$; $p = .03$. We chose not to merge PCs 1 and 2 into a single variable, because this would have produced a large drop in the extent to which the new PC reproduced the original acoustical variables. The original acoustical descriptors were strongly correlated with the respective PC, see Table 2. The first PC was labeled Hard_{SO}, because it included the most accurate specifier of sounding-object hardness, $\tan\varphi_{aud}$. The second PC was labeled Size_{SO}, because it contained the most accurate specifier of sounding-object size, F. The third PC contained the most accurate specifiers of hammer hardness, Lou_{att} and Lousl1, and hammer/sounding-object impact, SCG_{att}. The third PC was labeled Imp, because Experiment 2 had shown that in absence of feedback on response correctness, even trained participants were unable to focus on the most accurate hammer hardness information.

We computed two separate robust rank regression models for the data from each of the experimental conditions using the acoustical PCs as predictors. The dependent variable was the grand median of the hardness estimates across blocks of trials and participants, computed as for Experiment 2. The parameters of the regression models for data from all experimental conditions are presented in Table 5. In both regression models, the number of outliers was 7 out of the 27 data points. Both regression models explained a minimum of 96% of the variance in the hardness estimates and in their ranks.

Table 5
Summary of Multiple Regression Analysis of Hardness Estimates in the Population of Participants in Experiment 3

Variable	*B*	*SE* B	β	R_p^2
		Hammer condition		
Har_{SO}	−0.40	0.04	$-.47^{**}$	.85
Size_{SO}	−0.40	0.04	$-.47^{**}$	.85
Imp	−0.51	0.04	-0.60^{**}	.88
		Sounding-object condition		
Har_{SO}	−0.79	0.05	$-.87^{**}$	.95
Size_{SO}	−0.38	0.05	$-.42^{**}$	.70
Imp	−0.22	0.05	$-.25^{**}$	.44

Note. Hard = hardness; Imp = hammer/sounding-object impact; SO = sounding object; H = hammer; *B* = regression coefficient; β = standardized regression coefficient; R_p^2 = partial R^2. R_p^2 estimates the perceptual weight of the PC of acoustical descriptors.
** $p < .01$; $df = 16$.

Using the same procedure as for Experiments 1 and 2, we tested for significant differences in the perceptual weights of the acoustical PCs, both within conditions (six contrasts) and between experimental conditions (three contrasts). A family-wise Bonferroni correction was adopted, $\alpha = .0083$ and $.0167$, for within- and between-condition contrasts, respectively. The only significant within-condition contrast was observed for the SO condition, where $Hard_{SO}$ was weighted significantly more heavily than any other PC of acoustical descriptors, $p < .001$. For the H condition, although judgments appeared to focus on the PC related to the hammer/sounding-object impact, no significant difference emerged in the weighting of the three acoustical PCs, $p \geq .13$. Interestingly, the between-conditions contrasts revealed that the impact properties PC was weighted more heavily in the H condition, $p = .008$, and that the $Size_{SO}$ PC had the same perceptual weight in both conditions, $p = .15$. We also observed a tendency for the $Hard_{SO}$ PC to be weighted more heavily in the SO condition, although the significance level fell short of the critical α adjusted for multiple comparisons, $p = .04$.

Finally, we related the measures of perceptual weights of the acoustical PCs to the information accuracy scores computed on the database of sound sources (see Table 1). Figure 6 plots the perceptual weights of the acoustical PCs as a function of the median of the accuracy scores for the features they reduce. For the H condition, the perceptual weight slightly increased with information accuracy, although the differences between the perceptual weights of the different PCs were not significant. For the SO condition, a heavier weight was given to the most accurate PC, but the least accurate PC was weighted more heavily than the PC characterized by an intermediate accuracy. As emerging from the within-condition contrasts, the most accurate PC was weighted significantly more heavily than the two least accurate PCs, whose perceptual weights did not differ significantly. The same trends emerged when omitting from the analyses the accuracy scores for the best acoustical specifiers of hammer hardness, Lou_{att} and Lou_{s11}.

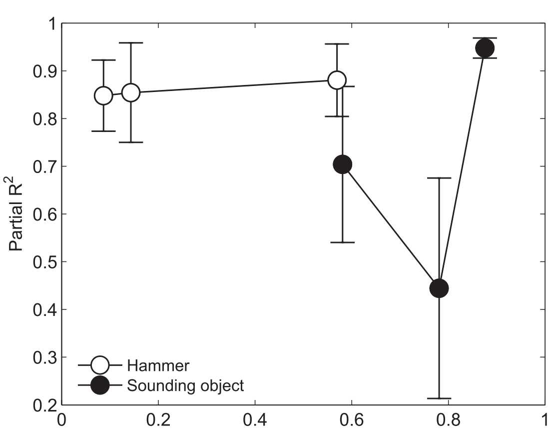

Figure 6. Perceptual weights of the principal components of acoustical descriptors in Experiment 3, as quantified by the R_p^2 measure of effect size, as a function of the median of the accuracy scores for the reduced acoustical features. Error-bars = ± 1 bootstrap standard error of the mean (10,000 bootstrap replicates).

Discussion

We investigated the estimation of the hardness of the hammer and sounding-object with a new group of untrained listeners. The stimuli were synthesized using a model that allowed the independent manipulation of the most accurate acoustical specifiers of the hardness and size of the sounding-object and of the hammer/sounding-object impact.

Consistently with results from Experiment 2, listeners integrated information over uncorrelated groups of acoustical features, independently of whether they judged the hammer or the sounding object. In particular, whereas hammer hardness ratings weighted equally information related to sounding-object hardness and size, on the one hand, and to impact properties, on the other, sounding-object hardness ratings weighted the accurate information for the sounding-object hardness more heavily. Thanks to the independent manipulation of sounding-object hardness and impact properties, we could test whether the results of Experiment 2 arose from a generalized bias towards focusing on the sounding-object material, independently of the judged object, or whether they resulted from a joint effect of information accuracy and exploitability. In line with the second hypothesis, sounding-object hardness information was less relevant to the perception of the hammer hardness, although the difference failed to reach statistical significance. Also, and more clearly, information related to the impact properties was weighted more heavily in ratings of hammer hardness. It is likely that similar perceptual criteria were used by the untrained participants in Experiment 2.

General Discussion

Empirical investigations on our ability to perceive nonspatial properties of the sound source almost invariably show that we do so by integrating information from multiple acoustical features. Despite our knowledge of what properties of a sound influence the estimation of given properties of the sound source, it is still unclear which principles determine what sound properties will dominate a specific perceptual judgment. We argued that the integration of information in source perception is governed by two independent factors. By drawing from ecological theories of human perception (e.g., Michaels & Carello, 1981), we hypothesized that the perceptual weight of an acoustical feature increases with the extent to which it accurately specifies the judged sound source. By drawing from theories of multisensory integration (e.g., Ernst & Banks, 2002), we hypothesized that the perceptual weight of an acoustical feature increases with the extent to which we are able to exploit the information it carries, as determined by discrimination, learning and memory abilities.

We conducted three behavioral experiments to verify our hypotheses. Participants estimated the hardness of two objects whose interaction generates an impact sound: a hammer and a sounding object. We derived quantitative measures of the accuracy of acoustical information from the analysis of a large database of impacted sound sources. Information accuracy was defined in statistical terms, focusing on the extent to which a sound property reliably discriminates the levels of a sound source property within the learning environment. Thus defined, information-accuracy scores were related to the perceptual weights of the same acoustical features in the behavioral tasks. Measures of information exploitability were inferred from measures of hardness discrimination performance (Experiment 1) and from measures of the extent to which participants in Experiment 1 benefited from training when asked to carry out a different hardness estimation task on a different set of stimuli, and in the absence of trial-by-trial feedback on performance (Experiment 2).

Consistently with previous studies on sound source perception, results from all experiments confirm that perceptual judgments integrate information from multiple acoustical features. Indeed, listeners did not focus exclusively on the most reliable acoustical information, as assumed by ecological theories of perception (e.g., Michaels & Carello, 1981), but integrated information from both accurate and less accurate sound properties. Consistently with the hypothesis of Kellman (1996), the adult listeners in the current study seem to have adopted a comprehensive and robust perceptual strategy, theoretically capable of operating even in the absence of the most accurate sources of sensory information about the environment.

In line with our hypothesis on the role of information accuracy, listeners focused on the most accurate information when receiving trial-by-trial feedback on performance (Experiment 1). However, in the absence of feedback (Experiment 2), the perceptual weight appeared to be modulated not only by the accuracy of acoustical information, but also by its exploitability. More specifically, the trained participants who estimated sounding-object hardness in Experiment 1 by focusing on accurate and easily exploited information continued to do so in Experiment 2. On the other hand, the participants who estimated hammer hardness in Experiment 1 by focusing on accurate, but less easily exploited, information, ceased to do so in the absence of trial-by-trial feedback. These participants focused instead on less accurate, but more easily exploited, information. These trends were confirmed in the last experiment where untrained listeners were presented with synthetic sounds. Here, sounding-object hardness was estimated by focusing on the most accurate information, which was easily exploited. Further, the most accurate information for the estimation of the hammer/sounding-object impact, which allowed less-than-perfect but higher-than-chance estimation of the hardness of both objects, was more heavily weighted in judging hammer hardness.

Overall, the results of this study point toward a joint influence of information accuracy and exploitability on the structure of the perceptual criteria. Thus, accurate information is generally more relevant perceptually, although accurate but not easily exploited information is perceptually secondary at best. When generalized to the perception of environmental sounds at large, the results of this study imply that the perceptual weight of the acoustical features can be fully predicted from two sets of measurements: firstly, task-dependent measures of the accuracy of the acoustical information within the environment in which source-perception criteria are acquired; secondly, task-independent measures of the ability of a listener to exploit the information carried by the acoustical features. A theory of source perception will benefit from further empirical tests of these predictions.

References

Agresti, A. (1996). *An introduction to categorical data analysis.* New York: Wiley.

Avanzini, F., & Rocchesso, D. (2001). Controlling material properties in physical models of sounding objects. In *Proceedings of the International Computer Music Conference 2001* (pp. 91–94). September, 2001, La Habana, Cuba.

Bachem, A. (1954). Time factors in relative and absolute pitch determination. *Journal of the Acoustical Society of America, 26,* 751–753.

Ballas, J. A. (1993). Common factors in the identification of an assortment of brief everyday sounds. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 250–267.

Ballas, J. A., & Mullins, T. (1991). Effects of context on the identification of everyday sounds. *Human Performance, 4,* 199–219.

Carello, C., Anderson, K. L., & Kunkler-Peck, A. J. (1998). Perception of object length by sound. *Psychological Science, 9,* 211–214.

Carello, C., Wagman, J. B., & Turvey, M. T. (2003). Acoustical specification of object properties. In J. Anderson & B. Anderson (Eds.), *Moving image theory: Ecological considerations.* Carbondale, IL: Southern Illinois University Press.

Chaigne, A., & Doutaut, V. (1997). Numerical simulations of xylophones. I. Time-domain modeling of the vibrating bars. *Journal of the Acoustical Society of America, 101,* 539–557.

Efron, B., & Tibishirani, R. J. (1993). *An introduction to the bootstrap.* New York, NY: Chapman & Hall.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415,* 429–433.

Ernst, M. O., & Bülthoff, H. H. (2004). Merging the senses into a robust percept. *Trends in Cognitive Sciences, 8,* 162–169.

Fletcher, N. H., & Rossing, T. D. (1991). *The physics of musical instruments.* New York, NY: Springer–Verlag.

Freed, D. J. (1990). Auditory correlates of perceived mallet hardness for a set of recorded percussive events. *Journal of the Acoustical Society of America, 87,* 311–322.

Fujisaki, W., & Kashino, M. (2002). The basic hearing abilities of absolute pitch possessors. *Acoustical Science and Technology, 23,* 77–83.

Gibson, E. J., & Pick, A. D. (2000). *An ecological approach to perceptual learning and development.* New York, NY: Oxford University Press.

Giordano, B. L. (2003). Material categorization and hardness scaling in real and synthetic impact sounds. In D. Rocchesso & F. Fontana (Eds.), *The sounding object* (pp. 73–93). Firenze, Italy: Mondo Estremo.

Giordano, B. L. (2005). *Sound source perception in impact sounds.* Unpublished doctoral dissertation, University of Padova, Italy.

Giordano, B. L., & McAdams, S. (2006). Material identification of real impact sounds: Effects of size variation in steel, glass, wood and plexiglass plates. *Journal of the Acoustical Society of America, 119,* 1171–1181.

Grassi, M. (2005). Do we hear size or sound? Balls dropped on plates. *Perception & Psychophysics, 67,* 274–284.

Grey, J. M., & Gordon, J. W. (1978). Perceptual effects of spectral modifications on musical timbres. *Journal of the Acoustical Society of America, 63,* 1493–1500.

Houix, O. (2003). *Categorisation auditive des sources sonores.* Unpublished doctoral dissertation, Université du Maine, France.

Hubert, M., Rousseeuw, P., & Branden, K. (2005). ROBPCA: A new approach to robust principal component analysis. *Technometrics, 47,* 64–79.

Iman, R., & Conover, W. (1979). The use of the rank transform in regression. *Technometrics, 21,* 499–509.

ISO 389–8. (2004). *Acoustics: reference zero for the calibration of audiometric equipment–Part 8: Reference equivalent threshold sound pressure levels for pure tones and circumaural earphones* (Technical Report). Geneva, Switzerland: International Organization for Standardization.

Kellman, P. J. (1996). The origins of object perception. In R. Gelman & T. Kit–Fong Au (Eds.), *Handbook of perception and cognition, Volume 8: Perceptual and cognitive development* (pp. 3–48). San Diego, CA: Academic Press.

Klatzky, R. L., Pai, D. K., & Krotkov, E. P. (2000). Perception of material from contact sounds. *Presence: Teleoperators and Virtual Environment, 9,* 399–410.

Kunkler-Peck, A. J., & Turvey, M. T. (2000). Hearing shape. *Journal of Experimental Psychology: Human Perception and Performance, 26,* 279–294.

Lakatos, S., McAdams, S., & Caussé, R. (1997). The representation of auditory source characteristics: Simple geometric form. *Perception & Psychophysics, 59,* 1180–1190.

Levitin, D. J., & Rogers, S. E. (2005). Absolute pitch: Perception, coding, and controversies. *Trends in Cognitive Sciences, 9,* 26–33.

Li, X., Logan, R. J., & Pastore, R. E. (1991). Perception of acoustic source characteristics: Walking sounds. *Journal of the Acoustical Society of America, 90,* 3036–3049.

Lutfi, R. A. (2001). Auditory detection of hollowness. *Journal of the Acoustical Society of America, 110,* 1010–1019.

Lutfi, R. A., & Liu, C. J. (2007). Individual differences in source identification from synthesized impact sounds. *Journal of the Acoustical Society of America, 122,* 1017–1028.

Lutfi, R. A., & Oh, E. L. (1997). Auditory discrimination of material changes in a struck–clamped bar. *Journal of the Acoustical Society of America, 102,* 3647–3656.

Martin, F. N., & Champlin, C. A. (2000). Reconsidering the limits of normal hearing. *Journal of the American Academy of Audiology, 11,* 64–66.

McAdams, S. (1993). Recognition of sound sources and events. In S. McAdams & E. Bigand (Eds.), *Thinking in sound: The cognitive psychology of human audition* (pp. 146–198). New York: Oxford University Press.

McAdams, S., Kudo, K., & Kirchner, H. (1998). Perceptual interaction of excitor and resonator properties in percussive instrument sounds. In *Proceedings of the 16th International Congress on Acoustics* (pp. 1885–1886). Seattle, WA: International Commission on Acoustics.

McAdams, S., Winsberg, S., Donnadieu, S., De Soete, G., & Krimphoff, J. (1995). Perceptual scaling of synthesized musical timbres: Common dimensions, specificities, and latent subject classes. *Psychological Research, 58,* 177–192.

Michaels, C. F., & Carello, C. (1981). *Direct perception.* Englewood Cliffs, NJ: Prentice–Hall.

Moore, B. (2003). *An introduction to the psychology of hearing* (5th ed.). San Diego, CA: Academic Press.

Mulaik, S. (2005). Variable selection. In B. Everitt & D. Howell (Eds.), *Encyclopedia of statistics in behavioral science.* Chichester, UK: John Wiley & Sons, Ltd.

Pastore, R. E., Flint, J. D., Gaston, J. R., & Solomon, M. J. (2008). Auditory event perception: The source-perception loop for posture in human gait. *Perception & Psychophysics, 70,* 13–29.

Repp, B. H. (1987). The sound of two hands clapping: An exploratory study. *Journal of the Acoustical Society of America, 81,* 1100–1109.

Rosenblum, L. D. (2004). Perceiving articulatory events: Lessons for an ecological psychoacoustics. In J. G. Neuhoff (Ed.), *Ecological psychoacoustics* (pp. 219–248). San Diego, CA: Elsevier Academic Press.

Roussarie, M. V. (1999). *Analyse perceptive de structures vibrantes simulees par modele physique [Perceptual analysis of vibrating structures simulated by physical modeling].* Unpublished doctoral dissertation, Université du Maine, France.

Rousseeuw, P., & van Driessen, K. (1999). A fast algorithm for the minimum covariance determinant estimator. *Technometrics, 41,* 212–223.

Ryan, T. P. (1997). *Modern regression methods.* New York: Wiley-Interscience.

Tucker, L. R. (1960). Intra-individual and inter-individual multidimensionality. In H. Gulliksen & S. Messick (Eds.), *Psychological scaling: Theory and applications* (pp. 155–167). New York: John Wiley & Sons.

Tucker, S., & Brown, G. J. (2003). Modelling the auditory perception of size, shape and material: Applications to the classification of transient sonar sounds. In *Proceedings of the 114th Convention of the Audio Engineering Society.* March, 2003, Amsterdam, the Netherlands.

Van Valkenburg, D., & Kubovy, M. (2004). From Gibson's fire to gestalts: A bridge–building theory of perceptual objecthood. In J. G. Neuhoff (Ed.), *Ecological psychoacoustics* (pp. 113–147). San Diego, CA: Elsevier Academic Press.

Verbeke, G., & Molenberghs, G. (2000). *Linear mixed models for longitudinal data.* New York: Springer-Verlag.

Verboven, S., & Hubert, M. (2005). LIBRA: A Matlab library for robust analysis. *Chemometrics and Intelligent Laboratory Systems, 75,* 127–136.

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.

Waterman, N. A., & Ashby, M. F. (1997). *The materials selector* (2nd ed.). London, UK: Chapman & Hall.

West, B. T., Welch, K. B., & Galecki, A. T. (2006). *Linear mixed models: A practical guide using statistical software.* Boca Raton, FL: Chapman & Hall/CRC Press.

Received September 19, 2007
Revision received July 28, 2009
Accepted August 6, 2009 ■