Auditory Perception of Geometry-Invariant Material Properties

Zhimin Ren, Hengchin Yeh, Roberta Klatzky, and Ming C. Lin

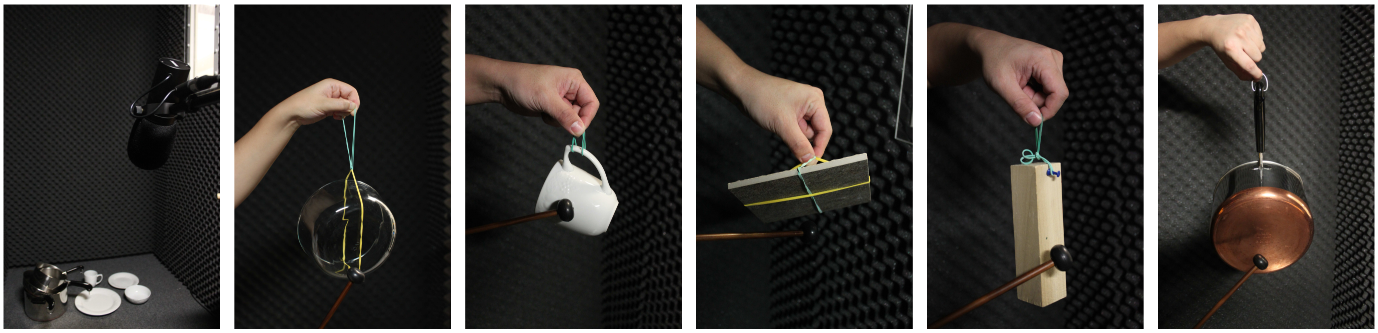

(a) Sound booth (b) Glass (c) Porcelain (d) Ceramic (e) Wood (f) Metal

Fig. 1: Recording setup: (a) the sound booth where recordings take place. Other figures (b) - (f) the setups for recording impact sounds from real-world materials: glass, porcelain, ceramic, wood, and metal, respectively.

Abstract— Accurately modeling the intrinsic material-dependent damping property for interactive sound rendering is a challenging problem. The *Rayleigh damping* model is commonly regarded as an adequate engineering model for interactive sound synthesis in virtual environment applications, but this assumption has never been rigorously analyzed. In this paper, we conduct a formal evaluation of this model. Our goal is to determine if auditory perception of material under Rayleigh damping assumption is "geometry-invariant", i.e. if this approximation model is transferable across different shapes and sizes. First, audio recordings of same-material objects in various shapes and sizes are analyzed to determine if they can be approximated by the Rayleigh damping model with a single set of parameters. Next, we design and conduct a series of psychoacoustic experiments, in subjects evaluate if audio clips synthesized using the Rayleigh damping model are from the same material, when we alter the material, shape, and size parameters. Through both quantitative and qualitative evaluation, we show that the acoustic properties of the Rayleigh damping model for a single material is generally preserved across different geometries of objects consisting of *homogeneous materials* and is therefore a suitable, geometry-invariant sound model. Our study results also show that consistent with prior *crossmodal* expectations, *visual* perception of geometry can affect the *auditory* perception of materials. These findings facilitate the wide adoption of Rayleigh damping for interactive auditory systems and enable reuse of material parameters under this approximation model across different shapes and sizes, without laborious per-object parameter tuning.

Index Terms—Sound synthesis, human perception of material.

1 Introduction

Realistic sound effects that closely correlate with visual stimulus play a vital role in many virtual environment (VE) systems and interactive 3D graphics applications, e.g. video games, immersive simulators, and special effects. With recent advances in high-quality audio generation, physically-based sound synthesis is gradually becoming a feasible and suitable approach for automatic incorporation of convincing sound effects in 3D graphics applications. These methods offer synthesized sounds based on material properties, object geometries, and physical contacts that excite the resonant objects.

Among various physically-based sound synthesis methods, *modal synthesis* [2, 22] is one of the most widely used real-time techniques in VE applications. It is highly efficient because it reduces complex vibrations of arbitrary geometries and materials to a linear combination of decoupled resonance modes. The geometry, characterized typically by shape and size, along with material parameters, determines the resonance modes obtained in the preprocessing step called *modal analysis*. When modeling resonant materials using modal sound synthesis, the *damping* component has always been a challenging issue, largely because the mechanism of energy dissipation for vibration is complex and not well understood. Moreover, modal decoupling is only feasible under certain damping models. *Rayleigh damping* [18] is one of the approximation models that enable such decoupling. As a result, it has been commonly adopted in rigid-body sound synthesis. However, to the best of our knowledge, though widely used in engineering applications, there has not been a formal analysis or rigorous evaluation of the Rayleigh damping model's transferability across different geometry (i.e. shapes and sizes). In other words, it is unknown if a single set of Rayleigh damping model parameters is sufficient for an arbitrary space of geometries or if the parameters would have to be "tuned" for changing geometry.

Without such an assumption, the Rayleigh damping model can only be applied on a per-object basis and a new set of damping parameters must be selected and tuned for every unique geometry – even *with the same materials*. This greatly limits the use of this approximation model and the adoption of modal sound synthesis in general, since finding appropriate Rayleigh damping parameters *per object* is usually non-trivial, tedious, and time-consuming. This process of ma-

- *Zhimin Ren is with UNC-Chapel Hill, E-mail: zren@cs.unc.edu.*
- *Hengchin Yeh is with UNC-Chapel Hill, E-mail: hyeh@cs.unc.edu.*
- *Roberta Klatzky is with Carnegie Mellon University, E-mail: klatzky@cmu.edu.*
- *Ming C. Lin is with UNC-Chapel Hill, E-mail: lin@cs.unc.edu.*

Manuscript received 13 September 2012; accepted 10 January 2013; posted online 16 March 2013; mailed on 1 May 2013.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

terial parameter tuning can quickly become prohibitively expensive for even a slightly complex VE scenario, where objects of different shapes with the same material are simulated. For example, the virtual fracture sound simulated by Zheng and James [28] is only feasible when assuming the same material parameters, including Rayleigh damping parameters, for the hundreds of fractured pieces.

In this paper, we examine the Rayleigh damping model's transferability across different shapes and sizes, using both real-world audio recordings and synthesized sounds to perform both objective and subjective analysis of this approximation model. Our goal is to determine if auditory perception of material under Rayleigh damping assumption is "geometry-invariant", i.e. if this model is transferable across different shapes and sizes. To achieve this goal, we have conducted an empirical analysis and a number of psychoacoustic studies in exploring human auditory perception of materials using the Rayleigh damping model across different geometric variations, as well as crossmodal perception of material under the influence of geometry.

The rest of the paper is organized as follows. In Sec. 2, we briefly describe the formulation of Rayleigh damping and related work on material perception in visual rendering and sound synthesis. Sec. 3 introduces our empirical study with real-world audio recordings. We analyze the recorded impact sounds of five sets of real objects. Each set contains several objects of the *same* material but *different* shapes or sizes. We verify if these recordings of the same material can be fitted to the same Rayleigh damping parameters with relatively small errors. Sec. 4 presents a psychoacoustic study to evaluate material similarity. Based on the responses from the subjects, we analyze the transferability of Rayleigh damping model with respect to variation in shapes and sizes. In Sec. 5 and 6, we discuss our findings, the application of these findings, limitations, and possible future directions of this work.

2 Background

2.1 Rayleigh Damping Model

Sound from rigid bodies is generated due to resonant objects' vibration. In order to model this process accurately and efficiently, *linear modal synthesis* methods [2, 22] are commonly adopted. It assumes small deformations during object vibration, thus its dynamics can be modeled as a linear system described by:

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{C}\dot{\mathbf{x}} + \mathbf{K}\mathbf{x} = \mathbf{f}, \tag{1}$$

where $\mathbf{x} \in \mathbb{R}^{3N}$ is the displacement vector of the system, and $\mathbf{M}$, $\mathbf{C}$, $\mathbf{K}$ represent the mass, damping, stiffness matrices, respectively. $\mathbf{M}$ and $\mathbf{K}$ can be acquired through finite element analysis [16], simple mass-spring formulation [17] and so on. In an *undamped* system, $\mathbf{M}$ and $\mathbf{K}$ can be diagonalized, and through generalized eigen-decomposition the solution of Eqn. 1 can be obtained, which is a series of *decoupled* harmonic oscillators, or *resonance modes*. Therefore, the complex dynamics of resonant objects are simplified and can now be computed efficiently. This process is called *modal analysis*, which is a standard structural analysis technique in engineering. However, if the *damping* term is present, the vibration dynamics can be reduced to a decoupled linear system only if $\mathbf{C}$ can be diagonalized as well as $\mathbf{M}$ and $\mathbf{K}$. Rayleigh [18] proposed a formulation for the damping matrix:

$$\mathbf{C} = \alpha\mathbf{M} + \beta\mathbf{K}, \tag{2}$$

which is a linear combination of mass and stiffness matrices, where α and β are Rayleigh damping coefficients. Given this simplification, solutions to the linear system in Eqn. 1 are:

$$q_i = a_i e^{-d_i t}\sin(\omega_i t + \theta_i). \tag{3}$$

In this equation, ω_i and d_i are respectively the angular frequency and the decay rate of the ith mode, while a_i and θ_i are the excited amplitude and initial phase determined by runtime excitation. We further observe that the Rayleigh damping assumption (Eqn. 2) and solutions to the dynamics formulation (Eqn. 3) define a frequency-decay relationship as a circle determined by Rayleigh damping coefficients α and β:

$$\omega_i{}^2 + \left(d_i - \frac{1}{\beta}\right)^2 = \left(\frac{1}{\beta}\sqrt{1-\alpha\beta}\right)^2. \tag{4}$$

This frequency-dependent decay rate model is a simplification of the complex mechanism of real internal material friction.

This simple damping formulation allows modal decoupling, and, therefore, it has been extensively used in rigid-body sound synthesis [4, 16, 17, 10, 20, 28, 19]. However, despite its extensive adoption, the Rayleigh damping model has never been formally evaluated for its transferability across varying shapes and sizes. In other words, it has not been formally studied and validated that the same set of Rayleigh damping coefficients along with the intrinsic material parameters, i.e. density and elasticity, preserve the same sense of material perception, if they are applied to objects made of the same materials but different shapes and sizes.

2.2 Related Work

Human hearing and auditory perception have been widely studied by researchers. Among them, Gaver [7] designed experiments to study the perception of everyday sounds, more particularly in sonic events, such as struck bars of wooden and metallic materials, and went on to apply his results to designing user interface with auditory icons. Wildes and Richards [27] studied recording audio of anelastic solids and determined that the *angle of internal friction*, $tan(\phi)$, is constant throughout all geometries of the same material. This work essentially defines a simple damping model, in which decay rate is linearly dependent on frequency. This damping model has been adopted by previous sound synthesis work (e.g. [6, 24]).

Klatzky et al. [11] designed perceptual experiments with synthetic sounds using the same damping formulation and studied the relationship between perceived resonant materials and the parameters in this sound synthesis model. In particular they found that the decay parameter τ_d, or equivalently, the internal friction coefficient $tan(\phi)$, is a better indicator than frequency alone in determining material similarity. This work suggests that the decay parameter can be used as a shape-invariant material property for synthesizing sounds. They also found that when subjects were asked to directly assign a gross material category for a given synthetic sound, it is the combination of both the frequency and decay parameter that determines their categorization.

However, the constant internal friction model is not sufficient. Krotkov et al. [12, 13] analyzed the recordings of hitting real world objects of different materials and observed that for a given material, the internal friction is not a constant but instead a function of frequency. They suggest that the shape invariance may be encoded in the functional form of the relation of $tan(\phi)$ and frequency, and proposed that a quadratic function appears to be a possible fit. In fact, the *Rayleigh damping* model is one such quadratic formulation for relationship between damping and frequency.

Giordano and McAdams [8] studied synthesized, impacted xylophone bars with varying material and geometric properties. In their physical model, two viscoelastic damping coefficients were used to describe a material, which is similar to Rayleigh damping. The relation of these properties to perceptual dissimilarity of the resulting sounds were studied, and a two-dimensional perceptual space was found to correlate with the material properties, namely the density and one of the two viscoelastic damping coefficients. Their result attests to the perceptual salience of energy-loss phenomena in sound source behavior. In another study, McAdams et al. [14] studied material categorization of recorded impact sound, and a large set of acoustic descriptors related to frequency, damping, and loudness. They found that a slightly modified measure, $tan(\phi_{aud})$, of damping is sufficient for recognition of gross material categories. For example, they combined steel with glass as a "gross" category steel-glass. They also combined wood and plexiglass, a special type of plastic, as plastic.

Multi-modal interaction in material perception involving both audio and visual was studied by Bonneel and Drettakis [3]. They varied the quality of synthesized sound and visual animation and studied subjects' material discrimination ability. Their study shows that high-quality audio rendering improves material perception, even when the visual rendering is low-quality. However, they did not show any correlation between visual and audio in material perception when virtual geometry vary. Visual perception of material reflectance is first

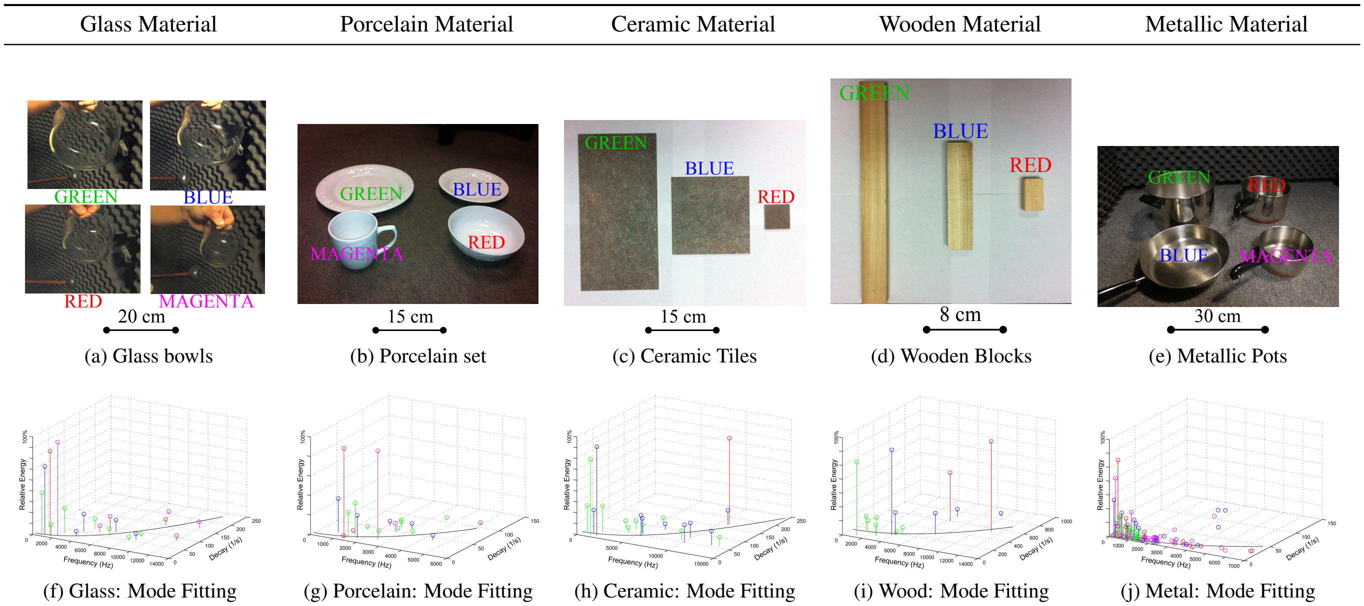

(a) Glass bowls (b) Porcelain set (c) Ceramic Tiles (d) Wooden Blocks (e) Metallic Pots

(f) Glass: Mode Fitting (g) Porcelain: Mode Fitting (h) Ceramic: Mode Fitting (i) Wood: Mode Fitting (j) Metal: Mode Fitting

Fig. 2: Fitting objects' resonance modes to the Rayleigh damping model. The top row shows real-world objects used in this experiment. The bottom row presents the fitting results, where the bottom plane represents the *frequency-decay* plane, the values in the height axis are relative energy, and the black curves on the frequency-decay plane visualize the fitted Rayleigh damping model. The color codes on the real objects match their extracted resonance modes in the same color.

studied through an exploratory psychophysical experiment in [26] to understand various influences on material discrimination in a realistic rendering setting. Their statistical analysis suggests that the accuracy of material perception is influenced by the geometrical shape of the object rendered with that particular material model. Nordahl et al. [15] synthesized footstep sounds in real-time for both solid and aggregate materials. They performed a perceptual study of floor material recognition for three groups of subjects. One group listened to real-world recorded footsteps, another group interactively generated footstep sounds by themselves and listened to the real-time synthesized sounds produced by the proposed system, and the third group listened to pre-recorded footstep sounds produced by the same system. Their study show that, in the *interactive* setup, subjects were able to identify synthesized floor materials at a comparable accuracy with real-world recordings, while the performance with pre-recorded sounds was significantly worse than the other two. This work provides interesting insights in how multi-modal interaction affects auditory material perception. However, visual elements are not included in this study.

3 Empirical Analysis of Real-World Recordings

In this experiment, we use recorded audio from real-world objects to evaluate the transferability of Rayleigh damping model across different geometry. To verify if the Rayleigh damping model is capable of capturing the intrinsic material damping that does not vary with the object's shape and size, we fit the recorded audio to a sound synthesis model using the Rayleigh damping assumption. If impact sounds from *same-material* objects in *different shapes and sizes* can be well approximated with the same Rayleigh damping model, this material model can be considered geometry-invariant across these objects. Five sets of real-world objects are selected for this experiment. Each set consists of three to four items made of the same material but with different geometry, i.e. varying shapes and/or sizes. The five sets are glass bowls (Fig. 2a), a set of porcelain dinnerware (Fig. 2b), ceramic tiles (Fig. 2c), wooden blocks (Fig. 2d), and metallic pots (Fig. 2e). The legend under these figures indicate the sizes of these experimental objects. In this section, we first describe the setup of our recording sessions. Then, We use an existing method to extract the *resonance modes* from the original recordings, and the summation of these key features accurately represent the recorded audio. Finally, we present the results for fitting these resonance modes to corresponding Rayleigh damping models. The fitting results are shown in the bottom row of Fig. 2, where the resonance modes' colors match the color codes of the real objects shown in the top row.

3.1 Recording Setup

Recordings were performed in a professional-quality sound booth, where all walls are padded with absorption materials to reduce reverberation effects, as shown in Fig. 1a. In order to generate impact sounds that best capture the intrinsic resonance properties of objects, we try to minimize their contacts with other articles. In most cases, rubber bands are used to suspend the object of interest, allowing the object to vibrate with minimum external damping due to contacts. The metallic pots are suspended by the attached metal loops (Fig. 1f). To reduce sounds coming from the striker during the impact motion, we adopt a mallet with a hard rubber head as the striking object. Special care is taken during the striking motion to minimize the swinging of the struck object, so that ringing sound effects are reduced. In order to limit the variation to only geometry and material, we manually control the striking motion's magnitude and direction to be as consistent as possible throughout all recordings. To diminish the hit point variation, all strikes are aimed at the center position of objects, for example the center point on the bottom of the glass bowls, metallic pots, and porcelain set. The recording setups for some examples are shown in Fig. 1.

3.2 Resonance Mode Extraction

Recorded audio is complex and high-dimensional data, which are difficult to directly map to any simple material model. As shown by Doel et al. [25] and Corbett et al. [5], many rigid-body impact sounds can be well approximated with the summation of a bank of *damped sinusoids* with different frequencies, decay rates, and amplitudes. Each damped sinusoid is considered one *resonance mode*, whose frequencies and decay rates are intrinsic to the particular object, while the amplitudes vary with the magnitude and location of an impact applied to the object. We adopt these *modes* as a high-level representation for the original sound.

We use the feature extraction method in [21] to determine the resonance mode representation of the recorded impact sound clips. This method uses an optimization framework that extracts modes from the original audio in a greedy fashion. *Power spectrograms* of the original

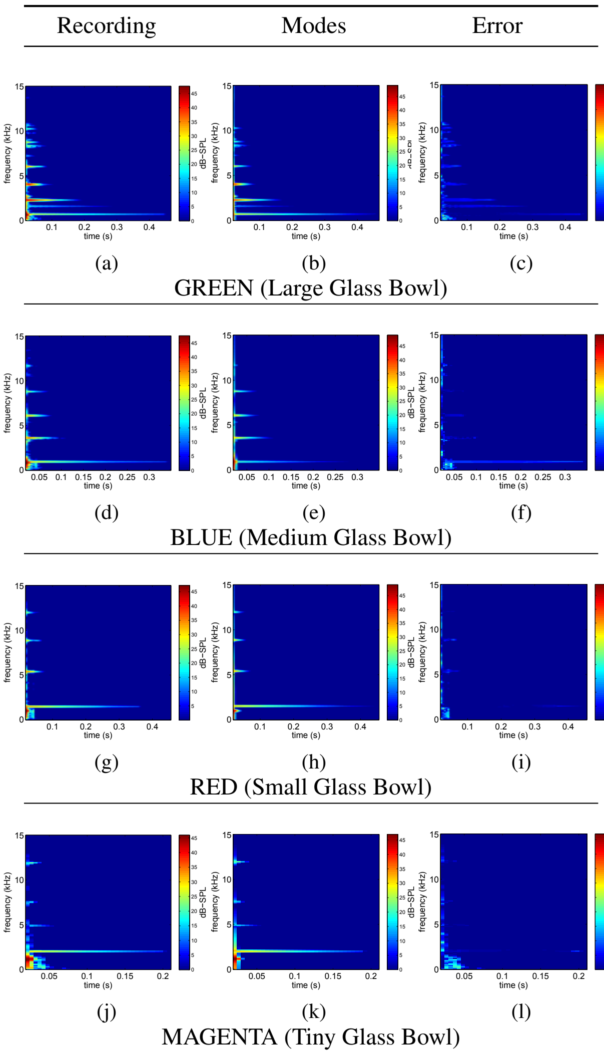

Fig. 3: Resonance mode extraction results for **glass** bowls of different shapes and sizes, as shown in Fig. 2a. Fig. 3a - Fig. 3l show, for each object, the power spectrograms of the recorded audio, the extracted resonance mode mixed audio, and the absolute error between the two. Resonance mode extraction results for other materials are shown in Appendix A.

recorded audio, the audio from mixing only the extracted resonances modes, and the absolute difference, i.e. error, between the two are shown for the glass bowls in Fig. 3, while the data for the experimental objects of other materials are included in Appendix. A (Fig. 10, 9, 11, and 12). The error plots show the extracted modes accurately capture the frequencies and decay rates of all prominent components in the recordings. Therefore, it is appropriate to use these modes to represent the original recordings for the purpose of studying the damping model, which defines a frequency-decay relationship for objects' vibration. For many objects, noticeable error appears in the range 0 - 1000Hz, which is quite possibly due to sound of the striker, i.e. the hard rubber ball, and the impact motion. How to separate the sound of striker and the struck object is still an open problem, which introduced error to the resonance mode analysis process.

3.3 Fitting Modes to the Rayleigh Damping Model

Once a resonance mode representation of a recording is acquired, we study how well the Rayleigh damping model can approximate these modes. We do so by fitting a curve following the Rayleigh damping model to these collected mode data points. For each material, the resonance modes of objects with different geometry are fitted to the same curve defined by one set of Rayleigh damping parameters. As shown in Fig. 2, we fit the curves on the 2D bottom plane to the observed (frequency, decay) pairs of modes. The values in the height axis represents relative energy of modes under a certain excitation. The relative energy values are used as weights in the least square regression, where the residual is defined as the difference between the observed and the predicted decay values. We weight the residual with relative energy because we want the fitted Rayleigh damping model to predict the more important modes (i.e. the higher energy modes) better than the less important ones. The fitting results for the five materials are shown as the black curves in Fig. 2f, 2g, 2h, 2i, and 2j. In Sec. 3.3, we statistically analyze the quality of the fit.

Quantitative Analysis of Goodness of Fit: In order to evaluate how well the curves fit the data, we compute the *coefficient of determination*, R^2, which is a widely used measure for assessing the *goodness* of regression using least squares techniques [23]. We adopted the standard weighted R^2 formulation,

$$R^2 = 1 - \frac{\Sigma w_i \times (y_i - \hat{y}_i)^2}{\Sigma w_i \times (y_i - \bar{y})^2}, \tag{5}$$

where $\{y_i\}$ are the decay values of the observed resonance modes, $\{\hat{y}_i\}$ are the decay values predicted by the Rayleigh damping model given the resonance modes' frequencies, $\bar{y}$ is the mean of $\{y_i\}$, i.e. the average value of observed decays, and $\{w_i\}$ are the weights, which are the relative energies of modes. Based on the standard interpretation of R^2 measure, an R^2 of 1 means the curve model perfectly fits the observed data, and the closer the value to 1 the better the fitting. The R^2 measures of the fitted Rayleigh damping models for the five materials in our experiment are listed in Table 1 ($p < 0.0001$ for all materials). This indicate the Rayleigh damping model generate predictions that are strongly and significantly correlated with the observed models of all materials. In four out of the five materials, the model accounts for approximately 75% of the observed variance in modes.

Table 1: Goodness of Fit for the Rayleigh Damping Model

	Glass Bowls	Porcelain Set	Ceramic Tiles	Wood Blocks	Metallic Pots
R^2 Measure	0.77	0.78	0.74	0.63	0.77

Notice the R^2 measure is noticeably lower for the wooden material compared with that of other materials. We believe that the anisotropy and other complex properties (e.g. heterogeneity of micro-structures) of the wooden material contribute to the fact that the simple Rayleigh damping model cannot fully reflect the damping phenomena of wood, hence the resonance modes fitted relatively poorly to the Rayleigh damping model. The relatively higher decay rates of the modes of wooden blocks may have also led to the poorer fitting. Nonetheless, the R^2 measures for all materials are reasonably high, indicating that in our experiment the Rayleigh damping approximation is accounting for a substantial, and highly significant, amount of the variance in the observed modes.

4 Perceptual Study on Material Similarity

In addition to the empirical experiment described in Sec. 3, we also conduct a psychoacoustic study where, in each trial, we ask subjects to determine if two sound clips played side-by-side are coming from objects made of the same *material*, while the objects can be of the same or different *geometry*. The study objective is to determine if the Rayleigh damping model can indeed capture the perceived material property sufficiently well to achieve transferability across different geometry.

Throughout this perceptual study, the independent variables are material and geometry (i.e. shape and size). The dependent variables that we measure as results are accuracy and confidence for experiments using recordings and consistency and confidence for those using synthetic sounds. Sec. 4.1 introduces what independent variables are used, and Sec. 4.2.1 describes how the study is designed to reasonably sample all independent variable combinations. Finally, Sec. 4.4 presents a detailed definition for the dependent variables and their values for the studies. We perform *within subject* study, where a single subject answer trial questions covering different combinations of independent variables, and in the end a *within subject* analysis of dependent variables is presented. In order to counterbalance, all the trial questions in this study appear in randomized order for every subject. In addition, the number of different-material synthetic sounds is very comparable to the number of same-material synthetic sounds, We also did not inform the subjects of the ratio of same material versus different material, and they go through the study treating each trial question as an independent incidence. Combining these factors, we believe our subjects do not have any assumption about the material identities before hand and are not biased to give a same-material or different-material answer in either way.

4.1 Audio Stimuli

In this experiment, subjects' perceived sense of materials is directly used as the indicator for determining whether Rayleigh damping model can be considered transferable across different geometry. However, human perception of materials is not solely dependent on the intrinsic material itself. It can also be affected by objects' geometry [11, 26]. We hope to study to what extent this effects the perception of real-world materials, and this finding serves as the baseline for interpreting the results from synthetic sound. Therefore, both *recorded* and *synthetic* audio clips are used as stimuli in our perceptual study.

For *recorded* audio stimuli, we use all the recordings acquired in Sec. 3. The first row in Fig. 2 shows pictures of the 18 objects for which impact sounds are used.

As to *synthetic* sounds, to explore the wide range of geometry and material variations, we selected a representative set of variations for generating the audio stimuli.

Shape variation: stick, cube, bunny, sphere, plate, and torus. They are shown in Fig. 4. These six sample shapes are chosen to represent shape variations such as complexity, dimensionality, and genus. For example, the simple cube shape is used, while the bunny shape is much more complicated. The plate is flat and circular, while the stick is much larger in one dimension than the other two. The sphere is a closed shape, while the torus has genus one. In addition, all shapes are solids that contain no cavity.

Size variation: small, medium, and large. We also vary the size of our sample shapes in order to study potential size-induced change in material perception. Three-size variations are adopted and illustrated on the example of bunny in Fig. 5. The smallest bunny is about 6*cm* tall, while the medium and large ones are respectively 2x and 4x the size of the small one. The same size variation is applied to all other shapes.

Material variation: metal, wood, glass, plastic, and porcelain. These five synthetic resonant materials are chosen to represent a variety of materials, and they are visualized on the sample shapes in Fig. 4.

In total, there are 90 variations arising from the combinations of the six shapes, three sizes, and five materials. Synthetic impact sounds for these 90 variations are generated using modal synthesis with the Rayleigh damping assumption, and they serve as the synthetic audio stimuli in our psychoacoustic experiments.

4.2 Study Design

In designing these experiments, we face two major challenges. Firstly, as described in the previous subsection, we have 18 recorded and 90 synthetic audio stimuli. If we aim to cover all variations in the stimulus space, picking two stimuli to form a question results in a huge number (nearly 12,000) of combinations which is infeasible for the study questionnaire. Secondly, human perception of material is inevitably affected by geometry variation. It is difficult to separate such effects in our study. Moreover, most people probably do not pay enough attention to auditory sensations in their daily lives to have closely observed the geometry effects in perceiving materials. Therefore, it is challenging to study auditory perception of materials across different geometry

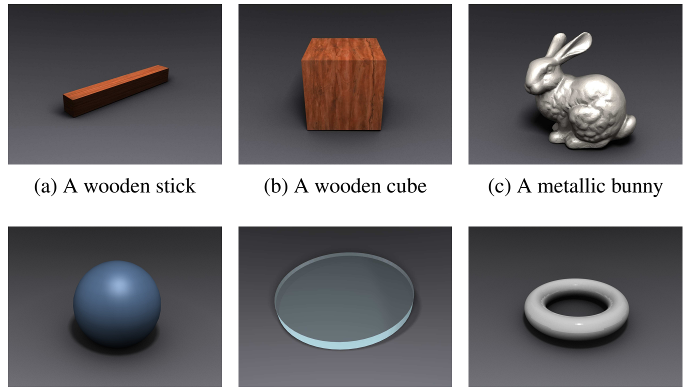

(a) A wooden stick (b) A wooden cube (c) A metallic bunny

(d) A plastic sphere (e) A glass plate (f) A porcelain torus

Fig. 4: Various **shapes** and **materials** used in the material similarity perceptual study described in Sec. 4. Six representative shapes: stick, cube, bunny, sphere, plate, and torus; five synthetic materials modeled with Rayleigh damping: metal, wood, glass, plastic, and porcelain.

Fig. 5: Three different **sizes** (1x, 2x, and 4x) for each shape, as shown on the metallic bunny example in this figure.

due to subjects' inability to distinguish variation in sound caused by geometry or material variation.

In this section, we first present an efficient stimulus sampling scheme that systematically picks pairing of audio stimuli to sample the combination space with a relatively small number of questions. Then, we describe our three-segment study procedure as an effort to better understand the perceived material variation due to geometric effects.

4.2.1 Stimulus Sampling

We randomly sample the complete stimulus combination space in the approach described below, where each subject is asked to complete a total of 56 trial questions. In particular, each subject judges six pairs from the recorded and 50 pairs from the synthetic stimulus set. We categorize our stimulus combinations based on their *material* and *geometry* configurations: identical or different material and identical or different geometry. This high-level grouping allows us to control the sample counts in each category and guarantees well distributed sample points that help us observe major trends.

Recorded stimulus sampling: Six trials are performed by each subject, and they are randomly selected from the 153 combinations made possible by picking any two from the 18 recorded stimuli. The random sampling follows the grouping listed in Table 2.

Table 2: Recorded stimulus sampling

	Material	Geometry	Count	
Group 1	Identical	Different	5	Total: 6 Trials
Group 2	Different	Different	1	

In Group 1, identical material and different geometry, one sample is selected for each real-world material out of the five we have, and the geometry combination is randomly selected. Group 2 is randomly

Table 3: Synthetic stimulus sampling

	Material	Geometry	Count	
Group 3	Identical	Different	30	Total: 50 Trials
Group 4	Different	Identical	10	
Group 5	Different	Different	10	

selected following the constraint of different material and different geometry. We pick more samples for Group 1 because we hope to gather more data with real-world *recordings* on how geometry affects material perception with the same material.

Synthetic stimulus sampling: Each subject is asked to complete a total of 50 trials for this category. The proposed sampling is outlined in Table 3.

Group 3 is the focus of this study, since it evaluates if the same sense of material is preserved across geometry variation when the synthetic stimuli are generated with the same Rayleigh damping material model. Geometry variation comes in two forms: shape and size. Therefore, Group 3 can be decomposed into three subgroups: different in both shape and size, only different in shape, and only different in size. 10 trials are performed respectively for each of these three subgroups. Particularly, the combination space is huge for the subgroup that is different in both shape and size. We propose the following scheme that achieves effective sampling for this subgroup. Fig. 6 illustrates the sampling scheme. First, 18 sample pairings are chosen from all shape-size combinations, and these samples satisfy that each chosen object is strictly selected twice in all combinations, and each pair is strictly different in both size and shape. Three such 18-combination groups are designed and color coded respectively in red, green, and black in Fig. 6. It appears these three groups evenly cover most combinations in the space. In each round of the study, one of the three groups is randomly selected, and 10 of the chosen group's 18 pairs are then randomly selected to represent the shape and size variation combination. Finally, for these 10 fixed geometry configurations, we randomly assign each of the five material choices to two of them. Thus, the 10 sample pairs for this subgroup are decided. For the subgroup of 10 pairs only different in shape, we fix the size configuration to be medium, randomly assign each of the five materials twice, and randomly combine shapes from the six options. Similarly, for the 10 pairs only different in size, we fix the shape configuration (five are fixed to be plates, and five are torus), two pairs for each material, and randomly combine sizes from the three options.

The 10 pairs in Group 4 and 5 follow the constraint of covering all possible material combinations (i.e. select any two out of five). For Group 4, an identical geometry configuration is randomly drawn for each pair, while for Group 5, two different geometry configurations are randomly selected for each trial.

With the above described sampling scheme, we define an approach that generates pairings in a random yet controlled fashion that provides us with experiments that cover a wide range of variants and focus on specific configurations that are central to our study (i.e. Group 3 in Table 3). Note that we did not include the group of identical material and identical geometry in either the recorded audio or the synthetic audio samplings. This is due to the subject's perfect identification rate for such pairings in our preliminary studies.

4.2.2 Study procedure

Our perceptual study is conducted in the format of online surveys. The interface of the study is shown in the accompanying video, and the study is designed to consist of the following three major parts, where each subject takes a 7-trial training session and then judges 67 stimulus pairs.

Training session: Geometry variation in objects leads to different qualities in sounds, which in turn affect subjects' auditory perception of material. It is challenging to separate the geometry and material influence in auditory perception. In order to take the geometry effect into consideration, our material similarity experiment includes a short training session, which shows subjects real-world sounds from objects in various materials and geometry. Subjects are firstly instructed to be aware that same-material objects can sound differently due to geometry variation. A video of impact sounds coming from four glass bowls that vary in shape and size are shown. Then they are asked to complete a seven-trial training, where each trial consists of two side-by-side audio clips. Subjects are asked to decide if the two clips are from the same material. Immediately after answering each training question, images of the resonance objects are revealed to subjects, which show audio and visual renderings of both the geometry and material of the experiment objects.

Material discrimination: The second part is an audio-only material discrimination study. Subjects are presented with two side-by-side audio clips and asked two questions for each trial. First, they are asked if the two audio clips come from objects made of *the same material*. Radio buttons for yes and no are provided for subjects to input their answers. Second, they are asked to rate how *confident* they are with their answer. Scores ranging from 0 to 10 represent 'not confident at all' to 'very confident'. The 56 trials sampled as described in Sec. 4.2.1 are conducted in this part of experiments.

Material discrimination with geometry visualization: The final part of this experiment is an audio-visual material discrimination study. The questionnaire is the same as the previous part, except that two side-by-side *images* corresponding to the two audio stimuli are also shown to subjects. These images are visual renderings of the resonance objects' *geometry*, and subjects are informed that they only carry geometry information and no texture or material clues. This allows us to explore how the added geometry visualization affects the subject's auditory perception of materials. A total of 11 trials are conducted, and they are sampled as shown in Table 4.

Table 4: Synthetic stimulus sampling with **geometry visualization**

	Material	Geometry	Count	
Group 6	Identical	Different	7	Total: 11 Trials
Group 7	Different	Identical	2	
Group 8	Different	Different	2	

A focus study: While the above study is relatively thorough, we hope to obtain more data in the category that is the most interesting to this study, i.e. Group 3 in Table 3, since we aim to evaluate if *Rayleigh damping* parameters transfer across geometry. Therefore, we also provide a focus perceptual study that only asks 21 trial questions after the 7-step training. 15 of the the 21 questions are subsamples chosen from Group 3 in Table 3, and the other 6 are randomly sampled in different material combinations, so that the study is more balanced with both same and different material pairs.

4.3 Participants

A total of 42 volunteer subjects, age between 21 and 45, were recruited for this perceptual study. 20 of them finished the full study, and among them 6 were female. The average age for this group is 28.40. The other 22 subjects completed the focus study, 8 are female, and the group average age is 32.27. All subjects reported normal hearing and performed the study at their own pace on a personal computer. All of them used headphones in the study for better audio quality, since frequency components at the high and low ends of audible spectrum might be inaudible through some consumer speakers.

4.4 Results and Analysis

The result of each trial is measured by the following two variables.

Consistency: Subjects are asked to answer if two audio clips are from the same material in each trial. For recorded stimuli, the concept of *accuracy* is directly adopted, since the ground truth of same or different materials for each trial pair can be determined, and an answer is correct or incorrect can be decided. For synthetic stimuli, we define *consistency*, which is analogous to *accuracy* for recorded stimuli. If subjects' answer is consistent with the material model assumption, we define consistency as 1.00. If not, it is defined to be 0.00. For example, if two audio clips in one trial are synthesized using the same material parameters, and the subject consider them the same material,

Shape/Size IDs:						
Shape:	1. bunny	2. cube	3. plate	4. sphere	5. stick	6. torus
size:	1. small	2. medium	3. large			

Color Codes:	
Group 1	Black
Group 2	Red
Group 3	Green

	(1, 1)	(1, 2)	(1, 3)	(2, 1)	(2, 2)	(2, 3)	(3, 1)	(3, 2)	(3, 3)	(4, 1)	(4, 2)	(4, 3)	(5, 1)	(5, 2)	(5, 3)	(6, 1)	(6, 2)	(6, 3)	row sum
(1, 1)																			0
(1, 2)																			0
(1, 3)																			0
(2, 1)		1																	1
(2, 2)	1		1																2
(2, 3)	1																		1
(3, 1)		1				1													2
(3, 2)																			0
(3, 3)	1			1															2
(4, 1)		1			1			1	1										4
(4, 2)	1		1	1		1	1												5
(4, 3)		1		1	1		1	1											5
(5, 1)			1					1	1		1								4
(5, 2)			1			1	1		1	1									5
(5, 3)	1	1		1	1			1											5
(6, 1)			1		1	1		1	1			1							6
(6, 2)	1		1			1				1			1		1				6
(6, 3)		1		1			1	1					1	1					6
column sum	6	6	6	5	4	5	4	6	4	2	1	1	2	1	1	0	0	0	
row + column sum	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	
sample sum	54																		

Fig. 6: Sampling schemes for subgroups: The number 1 marked in the spreadsheet cells indicates a selected combination. The shape and size IDs are listed in the top left corner of this table. The three different colors represent three different sampling subgroups. In each subgroup, each geometry of a distinctive size and shape is selected exactly twice, and all the combination pairs are different in both size and shape among the selected two geometry instances. Notice the combined three subgroups appear to randomly sample all possible pairings.

we assign 1.00 to the consistency of this trial. The mean consistency is essentially the proportion of subjects' answers consistent with the tested material model assumption.

Confidence: Besides the yes and no material discrimination question, subjects are also asked to rate their confidence with their decision. This 0 - 10 value indicates how confident the subject is, while 0 means not confident at all, and 10 is very confident. In other word, if the subject has difficulty or uncertainty in answering the material discrimination question, the confidence value of this trial should be low.

Results from the full-length studies are used in the analysis below. The focus study is solely designed to provide more samples for Group 3 in Table 3, so its results are only used in analyzing the across-shape and across-size cases described in Sec. 4.4.2. In all the results, we report both *means* and 95% *confidence intervals* (CI) for the accuracy/consistency and confidence values for each group, presented as CI centered around means. Where appropriate, *paired t-tests* [9] are performed to test the statistical significance of hypotheses on comparisons between two groups, and the *p-value*, which represents the probability of the observed result occurring by chance, is reported. We adopt .05 as the p-level for significance. The rest of this subsection includes the results and analysis of each data group. More observation and discussion are presented in Sec. 5.

4.4.1 Recorded stimulus trials

Table 5 shows the 95% confidence intervals of accuracy and confidence for all trials of *recorded* stimulus respectively in Group 1 and 2. Notice in Table 5, the accuracy rate is only 84.75% for Group 1. This indicates even with real-world recordings, when sounds of objects from identical materials are presented, subjects can be affected by the geometry variation and mistake identical materials as different. Perfect material discrimination across geometry variation is improbable. The variance in accuracy values is relatively large for Group 2, we suspect it is due to the small number of trials we performed for this particular category.

Table 5: 95% CI for accuracy and confidence for recorded audio trials

	Material	Geometry	Accuracy	Confidence
Group 1	Identical	Different	$84.75\% \pm 6.95\%$	7.46 ± 0.52
Group 2	Different	Different	$85.00\% \pm 16.06\%$	7.70 ± 0.77

4.4.2 Synthetic stimulus trials

Table 6 presents the 95% CI centered around mean consistency and confidence for each group throughout all synthetic audio trials. Notice the consistency rate for Group 3 is quite high, which indicates subjects perceive Rayleigh damping as transferable across geometries in a large proportion (around 76%) of the study trials. A paired two-tailed t-test between Group 3 and 5 indicates subjects are more capable of detecting mismatches than matches in material, when geometry differs ($t_{consistency}(20) = -3.01$, $p_{consistency} < 0.007$; $t_{confidence}(20) = -2.43$, $p_{confidence} < 0.025$). The same type of t-test between Group 4 and 5 fails to support the hypothesis that a geometric mismatch heightens reports of material mismatch ($t_{consistency}(20) = -1.94$, $p_{consistency} < 0.068$; $t_{confidence}(20) = -1.23$, $p_{confidence} < 0.233$).

Table 6: 95% CI for consistency and confidence for synthetic audio trials

	Material	Geometry	Consistency	Confidence
Group 3	Identical	Different	$76.73\% \pm 4.29\%$	7.25 ± 0.57
Group 4	Different	Identical	$81.17\% \pm 4.97\%$	7.65 ± 0.53
Group 5	Different	Different	$87.50\% \pm 4.69\%$	8.07 ± 0.46

In order to evaluate the differences among all materials, we also categorize the results based on the five materials in the study. Table 7 presents this result for each material throughout all synthetic audio trials. For all five materials, consistency and confidence are relatively high. Notice that the material of wood leads to the lowest performance.

Synthetic stimulus trials - same material, across shapes and sizes respectively: The focus of our study is to test if the same sense of material is preserved across different geometry (i.e. shapes and sizes), if the same material parameters including the same Rayleigh damping coefficients are assumed. Below, we present results categorized respectively into different shapes while sizes are fixed (Fig. 7) and different sizes while shapes are fixed (Fig. 8). All results in this part are calculated from the identical material trials in both the full-length and the focus study. Therefore, a total of 42 subjects' results are included. Fig. 7 shows results across all shapes: stick, cube, bunny, sphere, plate, and torus. Fig. 7a, 7b, 7c, 7d, and 7e present data for materials: wood, plastic, porcelain, metal, and glass, respectively. Fig. 8 shows results across all sizes: small, medium, and large. Fig. 8a, 8b, 8c, 8d, 8e present data for materials: wood, plastic, porcelain, metal, and glass, respectively. Once again, trials of wooden material yield one of the worst consistency rates. Additionally, the *small* objects in general appear to be identified as inconsistent with the material model more often than the other sizes. It also appears consistency varies more with shapes than sizes, which means, compared with sizes, a drastic shape change is more likely to lead subjects to identify sounds produced by the same material parameters as coming from different materials.

Table 7: 95% CI for consistency and confidence for synthetic materials

	Wood	Plastic	Porcelain	Metal	Glass
Consistency	$70.77\% \pm 7.92\%$	$90.81\% \pm 4.82\%$	$77.64\% \pm 7.23\%$	$81.69\% \pm 6.84\%$	$80.38\% \pm 7.49\%$
Confidence	7.26 ± 0.52	7.54 ± 0.62	7.25 ± 0.71	7.87 ± 0.60	7.24 ± 0.62

Synthetic stimulus trials - with geometry visualization: Table 8 shows results of the trials in which subjects are provided with visualization of the resonance objects' geometry. The consistency and confidence values are remarkably high. Group 7 has the highest consistency, and it is mainly due to that the geometry is identical. When a subject is shown the visualization of two identical geometries, it is clear the only variable is material. In this case, subjects can judge material similarity purely based on the variation in the perceived audio and not be affected by geometry variation at all. The results with geometry visualization are also categorized into different materials and shown in Table 9. The mean consistency and confidence values are generally larger than those of the audio only results (Table 7), while the standard deviations are also larger, which can be due to the smaller number of trials performed. In fact, in the comments left by several subjects, they specifically pointed out that the geometry visualization made the material discrimination task easier for them.

Table 8: 95% CI for consistency and confidence for synthetic audio trials with *geometry visualization*

	Material	Geometry	Consistency	Confidence
Group 6	Identical	Different	$87.14\% \pm 7.57\%$	7.62 ± 0.62
Group 7	Different	Identical	$95.00\% \pm 6.74\%$	8.53 ± 0.71
Group 8	Different	Different	$82.50\% \pm 12.87\%$	7.85 ± 0.74

5 Discussion

Through the empirical experiment in Sec. 3 and the perceptual study in Sec. 4, we make the following key observations.

The Rayleigh damping model can be considered geometry-invariant: In the empirical study, the Rayleigh damping model appeared to serve as a reasonably good approximation for five real-world resonance materials, based on the observed fitting results (i.e. R^2 measure in Table 1) for the experimental materials across different geometries. In addition, synthetic audio generated with the Rayleigh damping model were tested in our perceptual study. High consistency between these adopted synthetic materials and subjects' material discrimination were recorded (76.73% for Group 3 synthetic audio only trials in Table 6 and 86.14% for Group 6 synthetic audio with geometry visualization trials in Table 8). The consistency rates indicate that synthetic sounds of various geometry (i.e. sizes and shapes) using the same Rayleigh damping model are perceived as the same material at a very high percentage. In addition, the broken down across-shape and across-size consistency rates shown in Fig. 7 and Fig. 8 and the average consistency rates for each material recorded in Table 7 and Table 8 (with geometry visualization) are also relatively high, especially the ones with geometry information. Moreover, we need to consider that subjects are not capable of perfectly discriminating materials due to geometry variation. Evidence for this is shown in the recorded audio trials. Even if the underlying material is identical (no approximation with any model), subjects can mistake them for different materials. In fact, the mean consistency values for synthetic stimuli in Group 3 and 5 are not significantly smaller than those of recorded stimuli in Group 1 and 2, respectively. It suggests that synthetic stimuli with Rayleigh damping assumptions can be considered good approximations in terms of preserving the sense of materials that is comparable with that of real-world audio. From the results in our experiments, we verified when applying the same set of Rayleigh damping parameters across different geometries, the same sense of material is preserved to a large extent.

Multi-modal effects in auditory material perception: Respectively compare results in Table 6 and Table 8, and Table 7 and Table 9. We observe, with the added visualization of object geometry, subjects' material perception shows significantly higher agreement with the Rayleigh damping model. In fact, a paired two-tailed t-test between Group 3 and 6 has $t_{consistency}(20) = -3.34$, $P_{consistency} < 0.003$, and $t_{confidence}(20) = -2.29$, $P_{confidence} < 0.033$, and same type of t-test between Group 4 and 7 show $t_{consistency}(20) = -3.09$, $P_{consistency} < 0.006$, and $t_{confidence}(20) = -2.42$, $P_{confidence} < 0.026$. This strongly indicates that when visual geometry information is present, which is the case for most graphics and virtual environment applications, the Rayleigh damping model is perceived as geometry-invariant at an even higher rate. Therefore, Rayleigh damping assumptions should be readily adopted as a geometry-invariant material approximation model in most virtual environment applications. In scenarios, where multiple objects of various geometry are present, we can apply the same set of material parameters in Rayleigh damping model to them, and users would generally perceive them as bearing the same auditory material.

Rayleigh damping's limitations: Notice in Table 1, the fitting result is the poorest for the wooden blocks in this study. With synthetic audio samples (results in Table 7, Fig. 7 and Fig. 8), the wooden material also seems to be perceived as the least consistent with Rayleigh damping model. We posit that the Rayleigh damping model is not ideal for approximating anisotropic materials like wood, which display complex energy decay effects. In addition, the high decay rates of wood are possibly pushing the limits of Rayleigh damping assumption. Lastly, human's auditory perception in high frequency range is poor, and we believe it largely contributes to the worse agreement for smaller objects (as shown in Fig. 8), which generally have resonance modes of higher frequencies. The synthetic audio of porcelain and glass in our experiments also have higher frequencies, and their discrimination rates appear less consistent with the Rayleigh damping assumption (Fig. 7c and Fig. 8c, and Fig. 7e and Fig. 8e). Therefore, based on our studies, for relatively extreme cases like highly complex decay effects, large decay rates, and high frequency range, Rayleigh damping model is not ideal.

6 Conclusion and Future Work

In this paper, we have presented a number of experiments in which we examine the auditory perception of material across different geometry using the Rayleigh damping model for interactive sound synthesis in VR applications. We perform these studies both quantitatively and qualitatively by analyzing the real-world audio recordings and the synthetic sound clips generated by the Rayleigh damping model to determine if the material perception under this model is *geometry invariant*, i.e. does not vary across shapes and sizes.

Statistical analysis shows that the auditory perception of materials under the Rayleigh damping model for *homogeneous materials* is not influenced much by variation in shapes and/or sizes. However, our study results suggest that the Rayleigh damping model does not provide equally good approximation for materials with heterogeneous micro-structures, such as wood. Other more complex (perhaps more general but likely more compute-intensive) damping models [1] for capturing the material properties of sounding objects should be investigated and evaluated.

Reinforcing expectations based on well-known principles in cross-modal perception, our psychoacoustic experiments indicate that visual perception of geometry has noticeable effects on auditory perception of materials. This result is also consistent with study results in cross-modal perception and earlier study by [26] in which they found the visual perception of material is influenced by the geometry of objects.

These findings enable the wide adoption of Rayleigh damping in virtual environment applications for real-time modal sound synthesis and efficient reuse of material parameters under this approximation model across different geometry, thereby alleviating time-consuming per-object material parameter tuning.

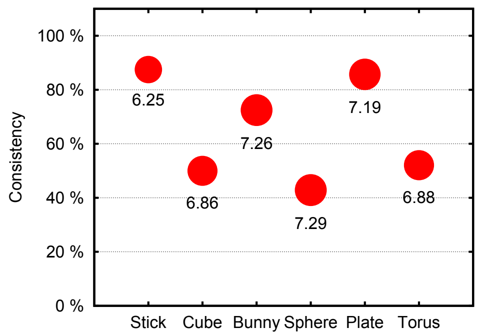

(a) Wood

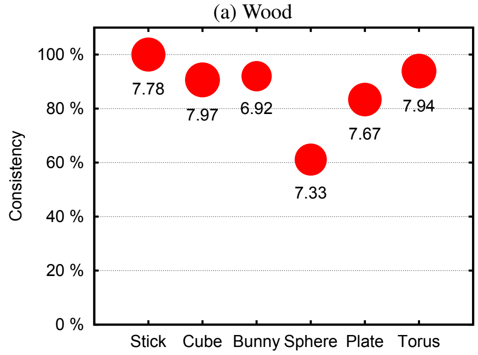

(b) Plastic

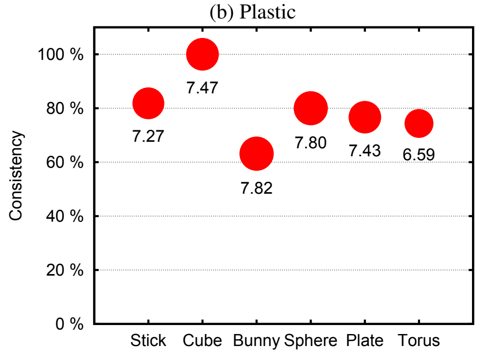

(c) Porcelain

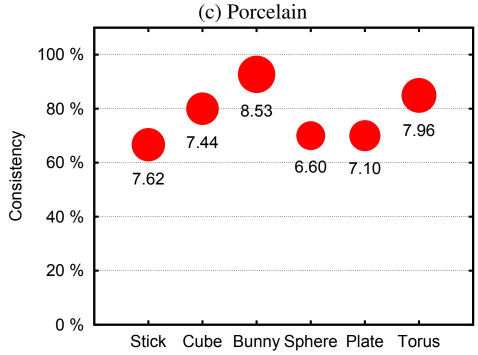

(d) Metal

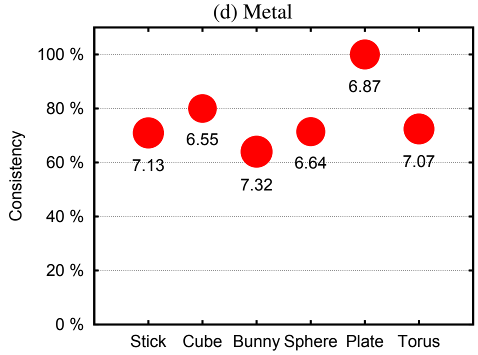

(e) Glass

Fig. 7: Consistency and confidence levels for synthetic audio trials across shapes for all materials. The radii of the disks represent confidence levels, which are also shown as the numbers below the disks.

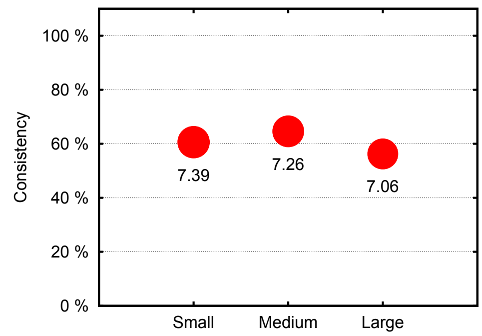

(a) Wood

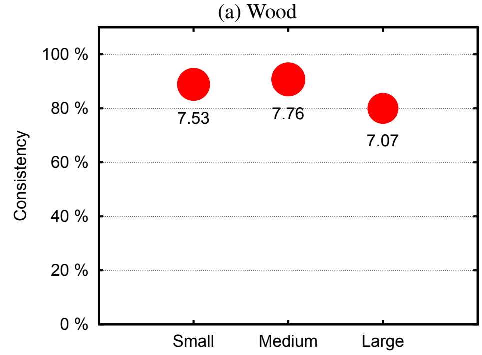

(b) Plastic

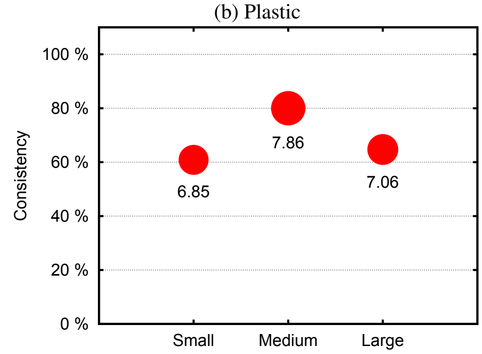

(c) Porcelain

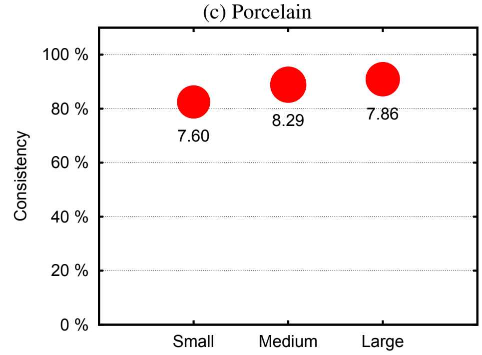

(d) Metal

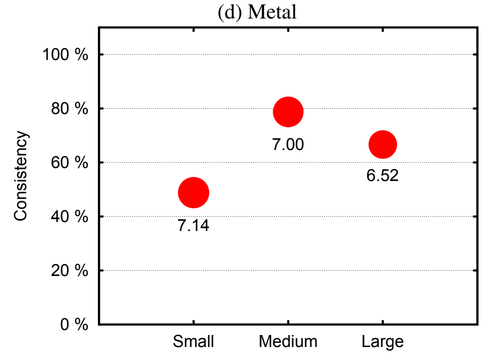

(e) Glass

Fig. 8: Consistency and confidence levels for synthetic audio trials across sizes for all materials. The radii of the disks represent confidence levels, which are also shown as the numbers below the disks.

Table 9: 95% CI for consistency and confidence for synthetic materials with geometry visualization

	Wood	Plastic	Porcelain	Metal	Glass
Consistency	$92.65\% \pm 7.72\%$	$92.11\% \pm 7.20\%$	$83.33\% \pm 13.60\%$	$87.72\% \pm 11.31\%$	$78.24\% \pm 15.89\%$
Confidence	7.78 ± 0.66	8.11 ± 0.69	7.53 ± 0.94	8.10 ± 0.63	7.49 ± 0.83

In the future, we hope to perform analytical and qualitative comparisons between the Rayleigh damping model and other damping models of higher degrees, as well as how different models affect sound synthesis algorithms both in rendered sound quality and computation costs. In addition, how to design perceptual studies to reduce the geometry variation effects in material discrimination tasks is worth studying. Perceptual studies on crossmodal (esp. auditory-visual) perception in virtual reality also demand more exploration.

Acknowledgments

This research was supported in part by ARO Contract W911NF-04-1-0088, NSF awards 0917040, 0904990, 100057 and 1117127, and Intel Corporation.

The authors wish to thank UNC GAMMA Research Group who helped providing feedback in the initial study, all the subjects who participated in the perceptual study, and the UNC Beasley Multimedia Center for providing the recording equipments.

References

[1] S. Adhikari and J. Woodhouse. Identification of damping: Part 1, viscous damping. *Journal of Sound and Vibration*, 243(1):43–61, 2001.

[2] J.-M. Adrien. Representations of musical signals. chapter The missing link: modal synthesis, pages 269–298. MIT Press, Cambridge, MA, USA, 1991.

[3] N. Bonneel, C. Suied, I. Viaud-Delmon, and G. Drettakis. Bimodal perception of audio-visual material properties for virtual environments. *ACM Transactions on Applied Perception (TAP)*, 7(1):1, 2010.

[4] P. R. Cook. *Real Sound Synthesis for Interactive Applications*. A. K. Peters, Ltd., Natick, MA, USA, 2002.

[5] R. Corbett, K. van den Doel, J. E. Lloyd, and W. Heidrich. Timbrefields: 3d interactive sound models for real-time audio. *Presence: Teleoperators and Virtual Environments*, 16(6):643–654, 2007.

[6] K. Doel and D. Pai. The sounds of physical shapes. *Presence*, 7(4):382–395, 1998.

[7] W. Gaver. *Everyday listening and auditory icons*. PhD thesis, University of California, San Diego, 1988.

[8] B. Giordano and S. Mcadams. Material identification of real impact sounds: Effects of size variation in steel, glass, wood, and plexiglass plates. *The Journal of the Acoustical Society of America*, 119:1171, 2006.

[9] D. Howell. *Statistical methods for psychology*. Wadsworth Pub Co, 2009.

[10] D. L. James, J. Barbič, and D. K. Pai. Precomputed acoustic transfer: output-sensitive, accurate sound generation for geometrically complex vibration sources. In *ACM SIGGRAPH 2006 Papers*, SIGGRAPH '06, pages 987–995, New York, NY, USA, 2006. ACM.

[11] R. Klatzky, D. Pai, and E. Krotkov. Perception of material from contact sounds. *Presence: Teleoperators & Virtual Environments*, 9(4):399–410, 2000.

[12] E. Krotkov, R. Klatzky, and N. Zumel. Analysis and synthesis of the sounds of impact based on shape-invariant properties of materials. In *Pattern Recognition, 1996., Proceedings of the 13th International Conference on*, volume 1, pages 115–119. IEEE, 1996.

[13] E. Krotkov, R. Klatzky, and N. Zumel. Robotic perception of material: Experiments with shape-invariant acoustic measures of material type. *Experimental Robotics IV*, pages 204–211, 1997.

[14] S. McAdams, A. Chaigne, and V. Roussarie. The psychomechanics of simulated sound sources: Material properties of impacted bars. *Journal of The Acoustical Society of America*, 115, 2004.

[15] R. Nordahl, S. Serafin, and L. Turchet. Sound synthesis and evaluation of interactive footsteps for virtual reality applications. In *Virtual Reality Conference (VR), 2010 IEEE*, pages 147 –153, march 2010.

[16] J. O'Brien, C. Shen, and C. Gatchalian. Synthesizing sounds from rigid-body simulations. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 175–181. ACM, 2002.

[17] N. Raghuvanshi and M. C. Lin. Interactive sound synthesis for large scale environments. In *Proceedings of the 2006 symposium on Interactive 3D graphics and games*, I3D '06, pages 101–108, New York, NY, USA, 2006. ACM.

[18] B. Rayleigh. *The theory of sound*, volume 2. Reprinted: Dover, New York, 1945.

[19] Z. Ren, R. Mehra, J. Coposky, and M. C. Lin. Tabletop ensemble: touch-enabled virtual percussion instruments. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, I3D '12, pages 7–14, New York, NY, USA, 2012. ACM.

[20] Z. Ren, H. Yeh, and M. Lin. Synthesizing Contact Sounds between textured Models. In *Virtual Reality Conference (VR), 2010 IEEE*, pages 139–146. IEEE, 2010.

[21] Z. Ren, H. Yeh, and M. C. Lin. Example-Guided Physically Based Modal Sound Synthesis. *ACM Transactions on Graphics*, 32(1), Jan. 2013.

[22] A. Shabana. *Vibration of discrete and continuous systems*. Springer Verlag, 1997.

[23] R. Steel and J. Torrie. *Principles and procedures of statistics: with special reference to the biological sciences*. McGraw-Hill Companies, 1960.

[24] T. Takala and J. Hahn. Sound rendering. In *ACM SIGGRAPH Computer Graphics*, volume 26, pages 211–220. ACM, 1992.

[25] K. van den Doel, P. Kry, and D. Pai. FoleyAutomatic: physically-based sound effects for interactive simulation and animation. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 537–544. ACM New York, NY, USA, 2001.

[26] P. Vangorp, J. Laurijssen, and P. Dutré. The influence of shape on the perception of material reflectance. In *ACM SIGGRAPH 2007 papers*, SIGGRAPH '07, New York, NY, USA, 2007. ACM.

[27] R. Wildes and W. Richards. Recovering material properties from sound. *Natural computation*, pages 356–363, 1988.

[28] C. Zheng and D. L. James. Rigid-body fracture sound with precomputed soundbanks. In *ACM SIGGRAPH 2010 papers*, SIGGRAPH '10, pages 69:1–69:13, New York, NY, USA, 2010. ACM.