To appear in ACM TOG 34(6) (SIGGRAPH Asia 2015).

Computational Design of Metallophone Contact Sounds

Gaurav Bharaj* David I.W. Levin† James Tompkin* Yun Fei‡ Hanspeter Pfister* Wojciech Matusik§ Changxi Zheng‡

* Harvard Paulson SEAS † Disney Research ‡ Columbia University § MIT CSAIL

Figure 1: Customized instrument. *A set of optimized 2D water jet cut animals and 3D printed cups form a musical scale. Provided with a user-supplied 3D shape and target sound characteristics of the desired frequencies and their amplitudes, our method optimizes the geometry of the object to realize the desired sound (see video).*

Abstract

Metallophones such as glockenspiels produce sounds in response to contact. Building these instruments is a complicated process, limiting their shapes to well-understood designs such as bars. We automatically optimize the shape of arbitrary 2D and 3D objects through deformation and perforation to produce sounds when struck which match user-supplied frequency and amplitude spectra. This optimization requires navigating a complex energy landscape, for which we develop *Latin Complement Sampling* to both speed up finding minima and provide probabilistic bounds on landscape exploration. Our method produces instruments which perform similarly to those that have been professionally-manufactured, while also expanding the scope of shape and sound that can be realized, e.g., single object chords. Furthermore, we can optimize sound spectra to create overtones and to dampen specific frequencies. Thus our technique allows even novices to design metallophones with unique sound and appearance.

CR Categories: G.1.6 [Optimization]: Constrained optimization H.5.5 [Sound and Music Computing]: Methodologies and techniques I.3.5 [Computational Geometry and Object Modeling]: Physically based modeling

Keywords: sound synthesis, shape optimization, computational fabrication, inverse problem

1 Introduction

Metallophone instruments produce sound directly by vibrating metal objects with a mallet strike. As such, the tone and timbre of this sound depend on the object shape and material. Over millennia, humans have developed design spaces of shapes which produce desired acoustic responses. However, discovering these spaces is expensive, time-consuming, and often requires trial and error. Even in cases where the design space is well known is additional expertise required. Take the glockenspiel: a set of suspended metal bars arranged as a keyboard which, when struck, produce a pure bell-like sound. Despite the fact that the size and shape of glockenspiel bars is well-studied, careful drilling of dimples on the underside of the bars is needed to tune the instrument. We explore whether we can use computation and digital fabrication to simplify the process of metallophone design.

Acoustic design has been heavily explored in mechanical engineering to avoid structural resonances, e.g., in beams or turbine blades. Algorithms attempt to optimize **either** the amplitude **or** the frequency response of an object, typically in response to an applied harmonic load. Such analysis has also been applied to musical instrument design to study the emission of sound from horns and the effect of violin bridges on tonal quality.

In contrast, we explore the problem of optimizing a given input geometry to control **both** the frequency **and** the amplitude of the vibrational response of metal objects to contact. We also exploit recent advances in automated manufacturing such as 3D printing, computer numerical control (CNC) milling, and water jet cutters to manufacture several novel struck metallophones which we compare to a professionally-manufactured glockenspiel. Simulation of contact sounds has long interested the graphics community, as has computational fabrication. We seek to bridge these two disciplines and explore how much control one can garner over the contact frequency spectra of complex geometries.

We tackle this problem using a new *functional specification* method which automates the design process and controls the acoustic response of a rigid object by optimizing its geometry. Given an initial parameterized shape and a sparse target frequency spectrum, we solve an inverse shape design problem using a global nonlinear optimization algorithm. During each step of the design space exploration, the current contact sound spectrum (computed using high-order finite element analysis) is compared to the user-desired frequencies and amplitudes. The search ceases once the two spectra are sufficiently close.

Significant challenges arise when trying to solve such optimization problems. The relationship between object geometry and intrinsic vibration modes is complex, and depends on both geometric features and material properties. As a result, we meet a constrained optimization problem that is high-dimensional and non-convex. We derive formulas for computing the derivatives of vibration modes with respect to the shape parameterization. These derivative formulas are general, allowing different shape parameterizations for different applications. Further, we develop

a new *Latin Complement Sampling* search algorithm which provides probabilistic bounds on design landscape coverage. This allows us to produce metallophones of sufficient quality (0–2% error) to serve as useful instruments with novel two- and three-dimensional shapes.

2 Related Work

A few works in graphics are closely related to our own. The example-based synthesis work of Ren et al. [2013] matches an object's sound spectrum to a recorded sound using modal analysis and a Nelder-Mead optimizer. This transfers recorded sounds to a virtual object by finding optimal material parameters. Throughout, the object geometry is given and unchanged.

Umetani et al. [2010] produced the first interactive design tool for *forward* metallophone fabrication. Our method solves the *inverse* problem non-interactively, affording three advantages: First, our method can optimize full 3D shapes as opposed to just 2D thin plates and produces pieces with richer timbres. Second, our method controls multiple frequencies, not just the fundamental (lowest) frequency of vibration, and our method controls the amplitude of each vibrational mode. This helps us create and dampen overtones (i.e., vibration frequencies higher than the fundamental, which is an important part of the timbre of the sound). Third, with no real-time requirement, we use high-order finite elements for accurate simulation, which reduces the need for post-fabrication correction.

Concurrently to this work, Hafner et al. [2015] optimized object thickness to control the lowest vibration frequency of the shape. They control only the smallest non-zero modal eigenvalue, but not overtones or vibration amplitude. In a sense, their problem is a simplified special case of the problem we tackle herein.

Acoustic Inverse Problem. Kac [1966] posed the isospectral shape question: "can one hear the shape of a drum?", or, do there exist two distinct shapes of membranes that resonate at the same frequencies? While answered "no" by Gordon et al. [1992a; 1992b], this famous question inspired many inverse acoustic works: given detected sound scattering patterns [Angell et al. 1997; Feijóo et al. 2004] or room echoes [Dokmanić et al. 2013], reconstruct the shape of the structures that affect sound propagation. Our work is also inspired by Kac's question, but has a very different problem formulation: the input specifies not only frequency values but also amplitudes, and our goal is to find a 3D shape composed of elastic materials. While not strictly an isospectral problem, these previous works give us insight into the difficulties we may encounter. Specifically, dense spectrum optimization of any shape would require a continuous space of isospectral shapes, something implied to be unlikely by Zelditch [2000]. We sidestep this by performing sparse spectrum control, leaving the rest of the frequency components uncontrolled but forcing their amplitudes to zero. This opens up a continuous space of shapes for our optimization method to efficiently explore.

Contact Sound Simulation. Computer music has modeled percussive instruments such as drums [Fontana and Rocchesso 1998] and xylophones [Essl and Cook 1999]. We exploit fast physics-based modal sound simulation from computer animation [van den Doel and Pai 1998; O'Brien et al. 2002; Raghuvanshi and Lin 2006; James et al. 2006; Chadwick et al. 2009; Zheng and James 2010; Zheng and James 2011; Lloyd et al. 2011; Ren et al. 2013] and acoustics [Chaigne and Doutaut 1997]. These rely on *modal analysis* [De Poli et al. 1991] (§4.1), which is widely used in computational mechanics. We extend these methods to compute derivatives of vibration modes from eigenanalysis, with respect to shape parameters, to use in our new optimization method.

Acoustics and Vibration in Engineering. Recent works have used computational design to produce musical geometries such as a saxophone or even speakers [Diegel 2013; Ishiguro and Poupyrev 2014]. Engineering has explored non-interactive inverse shape design for the vibrational modes or frequencies of an object. Some works [Yoo et al. 2006; Yu et al. 2010] focus on specific geometries (such as rotating beams or thin plates), while others are only concerned with the frequency [Choi and Kim 2006; Yamasaki et al. 2010] or the amplitude [Yu et al. 2013] of the vibration modes. Other works focus on objects under harmonic loads [Choi and Kim 2006] and are incompatible with the contact sound inverse shape design problem.

Some works optimize the geometry and topology of mechanical structures to control interaction with sound propagation, to reduce noise levels [Marburg 2002; Dühring et al. 2008; Barbieri and Barbieri 2006] or improve sound quality [Bängtsson et al. 2003; Wadbro and Berggren 2006]. These consider sound propagation via the wave equation or the Helmholtz equation. We focus on sound generation from modal vibrations of solid objects.

Our work is distinguished by its ability to optimize both sound frequency and amplitude for arbitrary volumetric geometry under a general external load. In addition, we consider how the object is supported, and optimize a stand to further suppress unwanted vibration modes. Finally, none of these works validate their results with fabricated objects as we do.

Functional Specification. Functional specification algorithms optimize object properties to meet a specific goal, and so require simulations to verify object design fitness. Recent methods have covered elastic deformation force, shape, and motion [Bickel et al. 2010; Skouras et al. 2012; Bickel et al. 2012; Zhu et al. 2012; Skouras et al. 2013], or appearance-based material distribution goals for subsurface scattering [Hašan et al. 2010; Dong et al. 2010], caustics [Papas et al. 2011] or reflectivity [Matusik et al. 2009; Weyrich et al. 2009].

Chen et al. [2013] abstract previous methods by goal, parameter reduction scheme, optimization method, and simulation algorithm. We contribute across the first three areas: sound spectra have been under explored in this functional specification space; our method allows different shape parameterizations during optimization; and we introduce a new stochastic-continuous optimization scheme specifically for our shape and sound problem.

Material Properties. Recent works find optimal material properties for deformation behavior under external forces for 3D fabrication [Bickel et al. 2009] and animation control [Lee and Lin 2012; Li et al. 2014; Xu et al. 2015]. Contact sounds are also affected by object material properties, where pleasing and repeatable sound requires stiffness for vibration. Changing metal properties within a single piece is difficult, therefore we fix material parameters and focus on optimizing geometric shapes for sound control.

Non-linear Optimization Methods. Our inverse shape design problem takes the form of a non-convex optimization. Sampling based schemes are commonly used to solve such problems [Hansen et al. 2003; Pettersson 2008; Bardenet and Kégl 2010; Snoek et al. 2012]. However, we propose a new shape optimization algorithm based on antithetical sampling [Nagaraj 2014], called Latin Compliment Sampling, which leverages local search (akin to Wampler et al. [2009]) and provides probabilistic bounds on search space coverage. We show that our method outperforms popular alternatives for complex optimization problems.

Contributions. Over existing works, we contribute:

- An algorithm to optimize the *entire* sound spectrum of a 3D object via shape variation. We use a new multi-objective for-

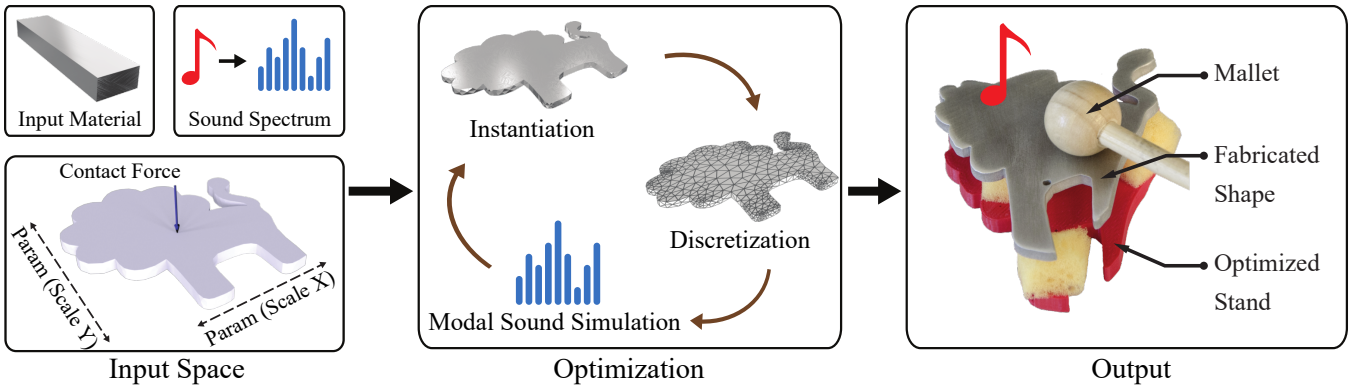

Figure 2: Overview. *Our algorithm for contact sound functional specification requires a shape, design parameters, material parameters and a frequency spectrum as inputs. We use a new global optimization technique, in conjunction with modal sound synthesis, to optimize shape parameters. Finally, we fabricate the computed shape (in this case, produced by a waterjet).*

mulation to optimize the frequency and amplitude of a struck linearly elastic object.

- *Latin Complement Sampling*: a hybrid search method tailored to non-convex constrained optimization, which provides probabilistic bounds on design landscape coverage.
- Analytical derivatives for sound frequency and sound amplitude with respect to arbitrary design parameters.
- A method for creating a stand for the fabricated object such that its sound quality is maximized.

3 Rational and Overview

Motivation. Acoustic design for instruments is concerned with frequency spectrum sparsity. For melodic instruments, most frequencies have negligible amplitude while a few are loud. These correspond to the note the instrument is to produce — its fundamental frequency — plus its overtones, which together describe the timbre of the sound. Producing a 'pure' note requires the amplitude of the fundamental frequency to be high compared to that of the overtones, while the remainder of the frequencies should be damped as much as possible. These observations inform the algorithm described below, which performs user-specified frequency spectrum sparsification.

Pipeline Overview. Our approach needs five user-supplied components (Fig. 2, left):

1. A shape parameterization and a design map to specific design instances — the 3D geometry.
2. A set of design parameters which specify the starting shape.
3. The object material parameters.
4. A contact region and force for striking the object.
5. The desired spectrum when struck, as a set of frequencies and corresponding amplitudes.

Our methodology assumes that the shape parametrization and material properties are known a priori. To begin, the user provides a sparse desired frequency spectrum and defines the contact region by labelling vertices on the surface of the initial design instance. Given this, we use isotropic scaling and spectrum simulation of a strike to match the fundamental frequency of our object to the desired fundamental frequency in a least squares sense. Then, for the overtones, we again simulate, and establish correspondences between the simulated frequency/amplitude values and their counterparts in the desired spectrum using absolute frequency difference. These correspondences are used to compute the amplitude and frequency cost functions that measure spectrum similarity (Fig. 2: Input). Our supplementary video provides an animation of this initialization process. During optimization we maintain this mapping by re-associating each user frequency/amplitude with the nearest eigenmode using absolute frequency and amplitude distance.

Our method produces a design instance that, when fabricated and struck as specified, produces a vibrational acoustic response with the desired frequencies and amplitudes. As our simulations assume vibration in free space — something that is impossible in the real world — we produce an optimized stand for each object which avoids damping user-desired frequencies and ensures the object can vibrate as needed (Fig. 2: Output).

We formulate this shape design task as a computational optimization problem (§4). The optimization variables are the shape parameters which describe a geometry instance in the design space. Fabrication limits such as object thickness and fabrication clearance impose constraints on these parameters. Complicating matters is the fact that the linear modal sound model (§4.1) imposes a highly nonlinear mapping between geometry shape and contact sound spectrum. This necessitates the development of a carefully-tailored solver which uses Complement Sampling to maximize its exploration of the design space. While searching the design space, we map the contact region from the initial design instance into world space using the shape parametrization, simulate contact spectra, and compared them to the user-supplied sparse goal spectrum. Once a design instance with sufficiently low error is located, the method terminates (Fig. 2: Optimization).

4 Problem Formulation

Formally, our mathematical derivation maps the shape design parameters to the resulting sound spectrum. To establish this map, we use the standard linear modal sound model [Shabana 1995; O'Brien et al. 2002; James et al. 2006; Cook et al. 2007]. This choice is motivated by the model's proven ability to accurately predict sounds of stiff materials such as metal.

We start by briefly reviewing the linear modal sound model that predicts sound spectrum from a provided shape (§4.1). Then, we formulate the core optimization problem (§4.2). This is solved using higher-order finite elements, for which we justify the need in the supplemental document.

4.1 Background on Modal Vibrational Sound

A linear modal sound model is built using a finite element approximation. Given the 3D geometry of a design instance and its material parameters, we discretize its volume using a tetrahedral mesh [Labelle and Shewchuk 2007]. Applying the finite element method leads to discrete equations of motion which govern the instance's dynamic response to external forces:

$$\mathsf{M}\ddot{\boldsymbol{d}} + \mathsf{D}\dot{\boldsymbol{d}} + \mathsf{K}\boldsymbol{d} = \boldsymbol{f}, \tag{1}$$

where M, K, and D are respectively the mass, stiffness, and damping matrices depending on the object shape and material; the vector $\boldsymbol{d} \in \mathbb{R}^{3n}$ describes the finite element nodal displacement with n nodes (i.e., n tetrahedral mesh vertices); and the right-hand-side vector $\boldsymbol{f} \in \mathbb{R}^{3n}$ stacks the external forces applied on tetrahedral nodes to excite the vibration.

D affects how fast the sound damps out. For stiff materials such as metal, the damping D is often small — this is the reason we hear long ringing sounds from a metallic tuning fork. Since D has little influence on the vibration frequencies, we assume it to be zero. Further, since in our formulation, we consider impact sounds, $\boldsymbol{f}$ is treated as a Dirac delta function which models an instantaneous impact force applied to a prescribed contact region.

The standard methodology for solving Equation (1) is *linear modal analysis* [Shabana 1995]. Provided the mass and stiffness matrices, we first solve a generalized eigenvalue problem:

$$\mathsf{KU} = \mathsf{MUS} \;\text{ and }\; \mathsf{U}^T\mathsf{MU} = \mathsf{I}. \tag{2}$$

This computes a modal shape matrix U and a diagonal eigenvalue matrix S. The former describes displacement patterns of individual vibration modes while the latter describes the square of the vibration frequencies. In other words, if we apply a force impulse $\boldsymbol{f}$ to an object, the resulting sound spectrum will consist of N individual frequency components, for which the frequency values ω_i and amplitudes a_i can be estimated using (Fig. 10):

$$\omega_i = \frac{1}{2\pi}\sqrt{\mathsf{S}_{i,i}} \;\text{ and }\; a_i = |\boldsymbol{f}^T\boldsymbol{u}_i|, \;\; i = 1...N, \tag{3}$$

where $\boldsymbol{u}_i$ is the i-th column vector of U.

Now we can summarize our *forward* use of the modal sound model. Given the shape of an instance, and a known fabrication material, we construct mass and stiffness matrices using a finite element mesh, and solve the generalized eigenvalue problem (Equation (2)). Then, using Equation (3), we predict the frequency spectrum produced when striking the instance at a user-specified location. This chain of operations establishes a relationship between the shape of the instance and its produced sound spectrum. It is this relationship that we exploit to solve the *inverse* optimal shape design problem.

4.2 General Contact Sound Optimization Problem

Notation. Let $\boldsymbol{p}$ denote the geometric parameters for shape design (i.e., the design-space parameters), and $\phi(\boldsymbol{p})$ denote a map between the design parameters and the derived 3D geometry (a design instance). The specific choice of $\phi(\boldsymbol{p})$ is application specific and depends on the type of deformations a user wishes to allow. For instance, for radially symmetric objects the parameters $\boldsymbol{p}$ are the 3D control points of a 1D curve and the ϕ is a rotational extrusion. To preserve this flexibility when presenting the problem definition here and in our optimization method (§5), we express $\phi(\boldsymbol{p})$ as a general map. We will explain specific parameterizations with our implementation details and experiments in §9. Further, we use $\omega_i(\phi(\boldsymbol{p}))$ and $a_i(\phi(\boldsymbol{p}))$ to denote the frequency and amplitude value of i-th component in the predicted sound spectrum, because as introduced in §4.1, both ω_i and a_i depend on the 3D geometry of the object via a generalized eigen-decomposition (Equation (3)).

Objective Functions. We wish to control ω_i and a_i where the frequency values control the pitch of the sound and the amplitudes control the loudness of specific frequencies. We consider this problem as a multi-objective optimization with two subgoals:

- **Frequency Composition.** We allow the user to select a subset of all frequency components and control their frequency values. Let $\mathcal{K}_f$ denote this subset, so $k \in \mathcal{K}_f$ is an index of a frequency component that the user wishes to control. For every $k \in \mathcal{K}_f$, the user specifies the desired frequency value ω_k^*. Then, we define an objective function:

$$E_\omega(\boldsymbol{p}) = \sum_{k\in\mathcal{K}_f} \frac{w_k}{\omega_k^*}\left[\omega_k(\phi(\boldsymbol{p})) - \omega_k^*\right]^2, \tag{4}$$

  where $w_k \in \mathbb{R}$ is a user-controlled weight assigned to balance the relative importance among those desired frequencies (e.g., the fundamental frequency against any overtones).

- **Frequency Amplitudes.** As the input to our system includes an expected contact region, we can compute the volume of each frequency component using Eq. (3). In practice, the user specifies a small region on the object surface. When computing a_i in Equation (3), we construct a force vector $\boldsymbol{f}$ that is uniformly distributed over the user-specified region. The direction of the force at each surface vertex is chosen to be normal to the surface and the summed magnitude of all forces is chosen to be one. Note that changing the force magnitude scales all frequency component amplitudes by a constant factor making the whole sound louder or softer, but not effecting the relative amplitudes of said components.

  We allow the user to select another subset, $\mathcal{K}_a$, of frequency components, for which the desired frequency amplitudes are specified. Here $\mathcal{K}_a$ and $\mathcal{K}_f$ can be independent (e.g., to dampen undesired frequencies). If $\mathcal{K}_a$ is identical to $\mathcal{K}_f$, then both the frequency values and amplitudes of those components are optimized. Let $a_k^*, \forall k \in \mathcal{K}_a$, denote the user-desired frequency amplitude. Then, the objective function takes the form:

$$E_a(\boldsymbol{p}) = \sum_{k\in\mathcal{K}_a} \frac{w_k}{\bar{a}_1}\left[a_k(\phi(\boldsymbol{p})) - a_k^*\right]^2, \tag{5}$$

  where w_k is again the weights to balance the relative importance of the amplitudes, and $\bar{a}_1$ is the amplitude of the fundamental frequency of the initial shape; used to normalize the frequency amplitudes across multiple components.

Multi-objective Optimization Problem. We combine both functions (4) and (5) into a multi-objective optimization problem and solve with a lexicographic approach [Branke et al. 2008]. Using multi-objective optimization allows us to explore the pareto frontier of optimal solutions without resorting to the non-intuitive weight twiddling associated with weighted sums of energy terms.

$$\boldsymbol{p}_\omega^* = \arg\min_{\boldsymbol{p}} E_\omega(\boldsymbol{p}), \tag{6}$$

where E_ω is defined in Eq. (4). Next, this result is used to initialize the amplitude control optimization which requires solving:

$$\boldsymbol{p}^* = \arg\min_{\boldsymbol{p}} E_a(\boldsymbol{p}), \;\; \text{s.t. } \omega_k(p) = \omega_k^*(\boldsymbol{p}_\omega^*) \;\; \forall k \in \mathcal{K}_f, \tag{7}$$

Algorithm 1 Latin Complement Sampling

1: Sample N parameter values $\mathcal{P}$ via Latin Hypercube Sampling
2: $t \leftarrow$ randomly choose one parameter sample
3: **for** $i \leftarrow 1, m$ **do** ▷ Repeat at most m times
4: Run local SQP starting from the parameter sample t
5: Cache all $\boldsymbol{p}$ values gradients reached by SQP
6: Reject unnecessary $\boldsymbol{p}$ values near local minima
7: Fit a GMM to the cached samples
8: $t \leftarrow$ the sample in $\mathcal{P}$ with the least GMM PDF value
9: **if** the GMM PDF of t is less than a threshold **then**
10: **return** ▷ optimization terminated
11: **end if**
12: **end for**

where E_α is defined in Equation (5). Meanwhile, the shape parameters $\boldsymbol{p}$ in both optimization problems need to satisfy certain fabricability constraints such as thickness and clearance limits. If there are N_L limits, we express them generally as:

$$C_i(\boldsymbol{p}) \geq 0, \quad i = 1 \cdots N_L, \tag{8}$$

and defer the discussion of their specific forms until §9.

Both objective functions (4) and (5) are highly nonlinear because of the complex dependence of ω_k and a_k on the 3D shape $\phi(\boldsymbol{p})$. There is no analytic expression for this dependence as it involves a high-dimensional generalized eigendecomposition (2). How can we tackle such an optimization problem numerically?

5 Optimization

To solve the above optimization problems, we propose a new optimization algorithm, Latin Complement Sampling (LCS), which better navigates complex energy landscapes and provides probabilistic bounds on its exploration (§5.2). To facilitate the computation, we further develop mathematical formulas for the derivatives of ω_i and a_i with respect to $\boldsymbol{p}$ (§5.3), so that we can directly compute their values with respect to a variety of parameterizations. We begin this section with a discussion of the rationale behind our algorithm before proceeding to the details.

5.1 Method Rationale

Observations. The particular form of our multi-objective cost function (Eq. (4) & (5)) depends on the specific choice of parameterization $\boldsymbol{p}$ and design map $\phi(\boldsymbol{p})$. However, in general, they share the following properties:

1. The dependence of our cost functions on $\boldsymbol{p}$ is continuous and differentiable, suggesting that a gradient-based approach could efficiently find *local* minima.
2. However, we have both linear and non-linear constraints. For instance, the constraints in Equation (7) are nonlinear.
3. Thus, the energy landscape is highly non-linear, with many local minima, cliffs from constraints, and large flat regions (Fig. 8). This suggests that *globally* a gradient-based approach would be easily confused.

These properties inform our choice of optimization scheme (see §8 for representative, low dimensional, example cost functions). Specifically, we seek a global optimization scheme that is compatible with the general constraints we encounter (Equation (7) and (8)); yet, we would like to exploit the differentiability of our energy functions to perform robust search for local minima.

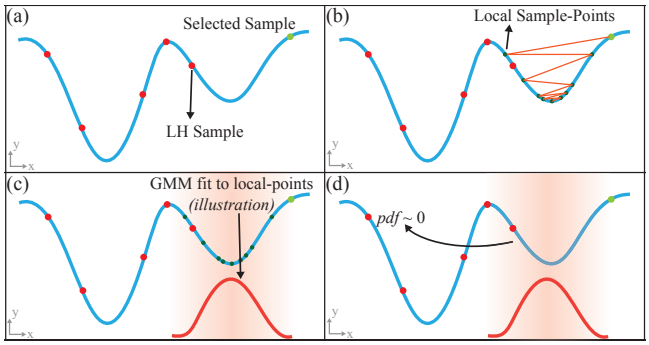

Figure 3: Latin Complement Sampling Steps. Top Left*: Choose sample.* Top Right*: Perform local search.* Bottom Left*: Fit GMM to local samples.* Bottom Right*: Sample from complement of GMM and repeat. (a) Latin Hypecube Sampling (b) Local Quasi-Newton Step (c) GMM Fitting (d) GMM Space-Complement.*

Previous Strategies. Global nonlinear optimization requires exploring large portions of the energy landscape, typically relying on sampling schemes. Algorithms like CMA-ES [Hansen et al. 2003] and Bayesian Optimization [Snoek et al. 2012] are good examples of this approach. In general, such methods attempt to fit a function to previous samples and estimate where new minima may appear. Exploration is then carried out near these minima [Pettersson 2008; Bardenet and Kégl 2010].

These algorithms have two issues in our context: First, it is difficult to include hard constraints in their formulations, and second they do not take advantage of available gradient information. Basin-CMA [Wampler and Popović 2009] attempts to fix these issues by performing a local Newton search at each sample point. This search serves as a projection operator onto the constraint set, as well as to move samples to local minima in the energy function. A Gaussian distribution is fit to the projected samples, then new samples are drawn from this distribution to continue exploring the energy landscape. However, as with previous works, it still relies on a sampling scheme which searches near anticipated minima.

In our case, this type of sampling is prone to failure as samples can fall into the first local minimum they find, leaving the remainder of the energy landscape unexplored (Fig. 6). Our solution to this problem is based on the antithetical sampling of Nagaraj [2014] but we extend the method with a Newton search procedure noting that *samples generated by a Newton search already denote the best solution locally, thus new samples should be generated far from them*. We call our variant of antithetical sampling *complement sampling* because it amounts to fitting a distribution to sets of local samples and then drawing samples from its complement (Fig. 3). We do this efficiently using Latin sampling and a sorting procedure to avoid explicitly computing the compliment of the sampling function (as in [Nagaraj 2014] for the purpose of analysis).

Method Overview. As outlined in Algorithm 1, our implementation of complement sampling, which we call Latin Complement Sampling (LCS), is divided into four phases:

1. Uniform sampling of the parameter space (Line 1 of Algorithm 1).
2. A local gradient-based minimization is performed from the chosen sample. The resulting locally-optimal parameter vector and its cost are cached (Line 4-5 of Algorithm 1).
3. Construction of a Gaussian mixture model using the cached parameter samples and their cost function values.

4. Selection of the next parameter sample using the above model (Line 7-8 of Algorithm 1).

The latter three steps are repeated until convergence or until a user-defined number of iterations is exceeded.

5.2 Latin Complement Sampling

Sampling Parameter Space. Suppose we are given a design-space parameterized by $\boldsymbol{p} \in \mathbb{R}^n$. These parameters need to satisfy box constraints $L_i \leq p_i \leq U_i, i = 1...n$ (see specific examples in Section 8), where p_i is the i-th component of $\boldsymbol{p}$.

Let $\mathcal{P} = \{\boldsymbol{p}_1, \boldsymbol{p}_2, ..., \boldsymbol{p}_N\}$ be samples from the parameter space. These samples will serve as candidates for performing local gradient-based optimizations. Therefore, ideally they need to explore the entire valid parameter space and sample uniformly. To this end, we construct $\mathcal{P}$ using Latin Hypecube Sampling (LHS) [McKay et al. 2000], a statistical approach known for its ability to spread sample points evenly across the sampled space.

Local Gradient Step. We begin by selecting a parameter sample $\boldsymbol{p}_s \in \mathcal{P}$, and using it as the initial point from which we perform a local gradient-based search (either on function (6) or (7) depending on the stage of the optimization). These are nonlinear least-squares problems with nonlinear constraints. Therefore, we adopt standard sequential quadratic programming (SQP) methods [Wright and Nocedal 1999].

The SQP method is itself an iterative algorithm, generating a new set of local samples $\mathcal{L} = \{\boldsymbol{l}_1, \boldsymbol{l}_2, ..., \boldsymbol{l}_M\}$ along a descent direction, terminating at a local minima. We cache samples $\boldsymbol{l} \in \mathcal{L}$ along with their associated cost functions values, filtering points that are too close together by Euclidean distance (threshold values range from 10^{-4} to 10^{-3}). These cached parameter points depict local regions that we have explored in the parameter space and we use them to inform our next choice of $\boldsymbol{p}_s \in \mathcal{P}$.

Complement Search. Throughout this iterative process, we need to continually choose parameter samples $\boldsymbol{p}_s \in \mathcal{P}$. During the first iteration, we select the sample randomly (Line 2 of Algorithm 1). After the first iteration, we choose the sampled parameters such that they maximize our exploration of the parameter space. Concretely, we fit an n-dimensional Gaussian mixture model (GMM) [Bishop et al. 2006] to the cached parameter samples from all previous iterations. Conceptually, given a parameter value $\boldsymbol{p}$, the resulting GMM indicates the probability that $\boldsymbol{p}$ was explored by previous gradient-based local navigation. Namely:

$$\text{prob}(\boldsymbol{p}) = \frac{1}{(2\pi)^{n/2}} \sum_{i=1}^{M} \frac{w_i}{|\Sigma_i|^{1/2}} \exp\left[-\frac{1}{2}(\boldsymbol{p}-\boldsymbol{\mu}_i)^T \Sigma_i^{-1} (\boldsymbol{p}-\boldsymbol{\mu}_i)\right].$$

Here, the number M of Gaussian components is determined by a KMeans++ clustering step [Arthur and Vassilvitskii 2007] to cluster the cached parameter samples; w_i, $\boldsymbol{\mu}_i$ and Σ_i are also learned from the sampled parameters [Dempster et al. 1977].

Next, we evaluate prob($\boldsymbol{p}$) for all $\boldsymbol{p} \in \mathcal{P}$ and find the parameter sample $\bar{\boldsymbol{p}}$ with the minimum GMM probability value. In other words, $\bar{\boldsymbol{p}}$ is the sample least likely to have been explored in previous iterations. Thus, we choose $\bar{\boldsymbol{p}}$ as the next starting point, $\boldsymbol{p}_s$, of a local SQP solve, and move on to the next iteration of the algorithm. A sufficiently large value of prob($\bar{\boldsymbol{p}}$) for all members of $\mathcal{P}$ (with a threshold of 0.3 in our examples) indicates that all the parameter samples have been covered, and so we terminate the algorithm (Line 9-10 of Alg. 1). The GMM model also provides us a probabilistic bound for the optimality of the given energy.

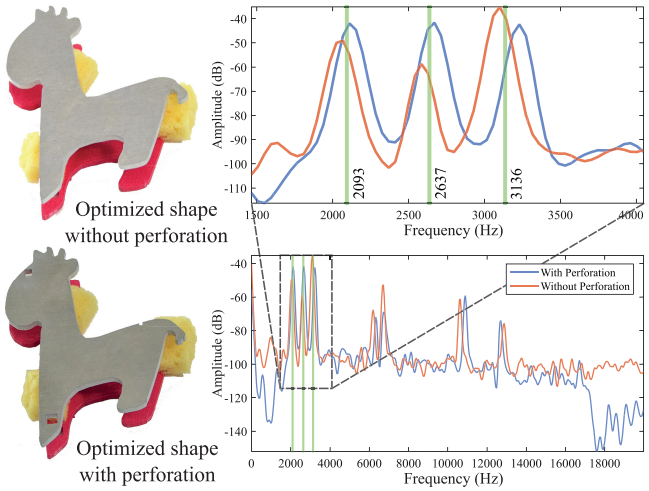

Figure 4: *Top left: Shape optimiziation alone (no perforations) can produce a zoolophone key with desired target frequencies (red curve) but with limited control of their amplitudes. Bottom left: Additional perforation optimization allows for control of frequency amplitudes (blue curve).*

5.3 Cost Function Gradients

One piece of our algorithm that remains to be presented is the computation of cost function gradients with respect to the design-space parameter $\boldsymbol{p}$. We could use finite differencing to compute gradients; however, this would require $2n$ cost function evaluations where $n = |\boldsymbol{p}|$ is the number of design-space parameters. Since each cost function evaluation requires solving a generalized eigenvalue problem, this approach can be slow. Worse still, finite differences are known for suffering from inaccuracy when functions vary rapidly. As illustrated in Figure 8, such rapidly changing functions are indeed what we encounter.

Using the chain rule we can compute gradients of both the modal sound frequency and amplitude cost functions, with respect to $\boldsymbol{p}$. The most challenging step is the computation of $\frac{\partial \omega_i}{\partial p_j}$ and $\frac{\partial a_i}{\partial p_j}$. We defer the detailed derivations to supplemental material, and present the final formulas:

$$\frac{\partial \omega_i}{\partial p_j} = \frac{1}{4\pi\sqrt{\mathsf{S}_{i,i}}} \boldsymbol{u}_i^T \left(\frac{\partial \mathsf{K}}{\partial p_j} - \mathsf{S}_{i,i} \frac{\partial \mathsf{M}}{\partial p_j}\right) \boldsymbol{u}_i \text{ and}$$

$$\frac{\partial a_i}{\partial p_j} = -\boldsymbol{f}^T (\mathsf{K} - \mathsf{S}_{i,i}\mathsf{M})^+ \left(\frac{\partial \mathsf{K}}{\partial p_j} - \mathsf{S}_{i,i} \frac{\partial \mathsf{M}}{\partial p_j}\right) \boldsymbol{u}_i + \frac{1}{2} \left(\boldsymbol{u}_i^T \frac{\partial \mathsf{M}}{\partial p_j} \boldsymbol{u}_i\right) \boldsymbol{f}^T \boldsymbol{u}_i,$$

where, following the notation in §4.1, K and M are respectively the stiffness and mass matrices; $\mathsf{S}_{i,i}$ is the i-th eigenvalue; and $\boldsymbol{u}_i$ is the i-th column vector of the modal matrix U. $(\mathsf{K} - \mathsf{S}_{i,i}\mathsf{M})^+$ is the pseudo-inverse of $\mathsf{K} - \mathsf{S}_{i,i}\mathsf{M}$.

6 Amplitude Control via Perforation

We observe experimentally that cutting small holes in an object has a negligible effect on its frequency response but a large effect on the amplitude of its vibration modes. This observation is of practical value as, in certain cases, it allows us to remove the constraints from our amplitude modulation cost function (Equation (7)), significantly simplifying it. We implement this by allowing $\phi(\boldsymbol{p})$ to encode the position and edge length of small square perforations on the object surface. We draw inspiration

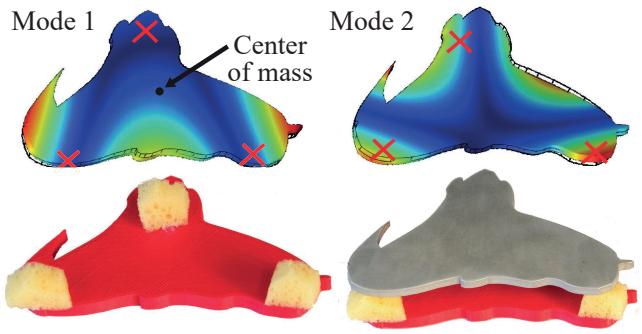

Figure 5: Stand creation process. Top: *The first two modal shapes inform where the supports must go (red crosses). Here, Mode 1 is desirable, while Mode 2 is not.* Bottom: *3D printed stand with foam-spikes for support and dampening.*

from small indents drilled into glockenspiel keys to damp spurious vibrations. Because indents are difficult to drill accurately, we chose to use perforations (which can be cut or included in a 3D print) instead. Figure 4 shows how perforations can be used to adjust the amplitude of peaks in an object's frequency spectrum.

7 Free Vibration with Stands

Our physical model assumes that the object itself is vibrating in free space, a condition impossible to replicate in the real-world. Anyone who has ever watched a drummer damp a ringing cymbal with his hand knows that contact can greatly affect the qualities of sound due to vibration. To ensure optimal sound production from our manufactured pieces, we construct optimized stands to avoid damping as much as possible. This is comparable to the rubber, string, or felt rail underneath the nodal point of a glockenspiel bar. Our stand creation methodology has three steps: determining candidate supporting vertices, sorting vertices based on their potential to induce damping in desired or undesired frequencies, and selecting a concrete subset of these vertices to support the object.

We rely on the observation that *a frequency is damped if its modal shape is not allowed to vibrate as if it were free*. This tells us that ideal support vertices have small maximum displacements in all desired frequencies and large minimum displacements in all undesired frequencies. In our case, we find a reasonable set of support vertices using an efficient 1D search over two sorted lists of vertices: one which contains vertices sorted in ascending order of their normalized maximum displacement in user desired frequencies, and a second which contains vertices sorted in descending order of their normalized minimum displacement in undesired frequencies. We choose the smallest subset of vertices from the beginning of these two lists that make our metallophone statically stable. Please see our supplemental material for complete details of the algorithm.

Stand Shape Creation. Once a supporting vertex set has been chosen, we create geometry for the stand using the CSG union of a flattened cuboid (the base) and upward pointing thin half-ellipses centered on each vertex. An example 3D-printed plastic stand is shown in Figure 5. The supports are covered with foam padding to avoid dampening from direct contact with the plastic. Stands are manufactured in the shape of the supported geometry to ease alignment and so clarify how each object is to be positioned.

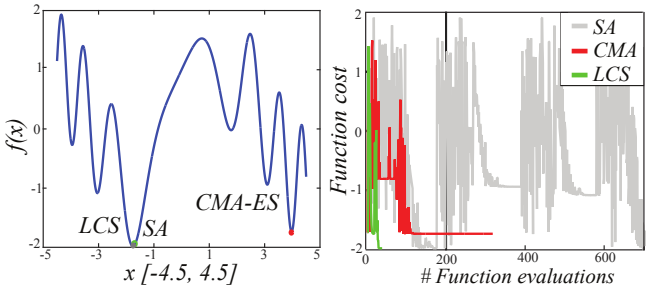

Figure 6: *1D function example comparing LCS to SA and CMA-ES.* Top: *Function and found minima.* Bottom: *Energy vs. function evaluation count.*

8 Validating Latin Complement-Sampling

We compare our LCS algorithm to two methods commonly used in computer graphics for stochastic optimization and computational fabrication: Simulated Annealing (SA) [Kirkpatrick et al. 1983] and Covariance Matrix Adaptation (CMA-ES) [Hansen et al. 2003] (see [Chen et al. 2013] for a review). To show the benefit of LCS, we also compare to Random Sampling (RSM) followed by local Newton's search from each sample.

First, we examine the behavior of LCS, CMA-ES and Simulated Annealing on a 1D non-convex cost function given by $f(x) = \sin(x) + \cos(x^2)$, over the interval $-4.5 \leq x \leq 4.5$ (Fig. 6(a)). We initialize all algorithms with the same parameter values and report the number of iterations required for each to converge to a minimum. Simulated Annealing is able to reach the global minimum in 700 iterations/function-evaluations while CMA-ES fails to find the same minimum as its initial samples fall into a nearby local minimum and are unable to escape. In contrast, LCS converges to the global minimum in 19 iterations (Fig. 6(b)).

Next, we show the behavior of all four algorithms on two well-known non-convex benchmark problems from the optimization literature: the Egg function and the Holder function (Fig. 7). We ran each algorithm 50 times, with random start points, and report the mean number of iterations until convergence to the global minimum. CMA did not converge (DNC) for either test while Simulated Annealing failed to converge on the Egg Function. LCS and RSM converged to the global minimum in both cases. However, LCS requires on average half the iterations of RSM, showing that LCS explores the parameter space more efficiently.

2D Scaling Frequency Optimization. We explore the performance of each algorithm on a sound fabrication test. We perform a frequency optimization task (Eq. (4)) using an extruded toric geometry as input (Fig. 8). The goal of the task is to compute a geometry instance with a fundamental frequency of 261.6 Hz (Middle C on the piano keyboard). Our design parameters, $\boldsymbol{p}$, control scaling in the x and y directions and, are limited to $0.5 \leq x, y \leq 1.5$ with simple box constraints (Fig. 8(a)). This produces a non-linear energy function landscape for this simple case (Fig. 8(b)). As above, we seed all algorithms with identical starting points. In this case, neither SA nor CMA-ES reach the global minimum while LCS finds a lower cost solution in fewer iterations. This is important because each function evaluation requires the solution of an expensive generalized eigenvalue problem (Tab. 1).

9 Metallophone Results

We use our parameterization-independent acoustic optimization to design 2D and 3D struck metallophones. For each result, we describe the input goal, the details of the parameterization used, and

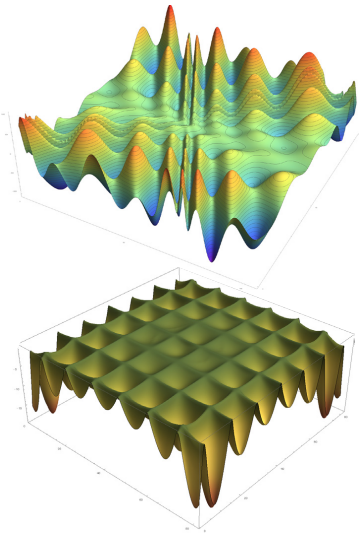

Egg Function

Method	Mean Iterations
SA	DNC
CMA	DNC
RSM	81
LCS	37

Holder Function

Method	Mean Iterations
SA	1866
CMA	DNC
RSM	466
LCS	266

Figure 7: *A comparison of Simulated Annealing (SA), Covariance Matrix Adaptation (CMA), Random Sampling + Local Search (RSM), and Latin Compliment Sampling (LCS) using the well-known Egg and Holder function benchmark problems. Each algorithm was run 50 times with random start points. We record the mean number of iterations required to converge to the global minima. DNC denotes that the algorithm did not converge to the global minimum.*

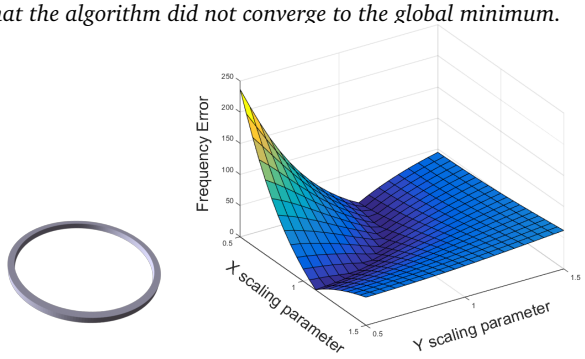

Figure 8: Left: *Toric section free to scale in x and y.* Right: *Energy over scale space showing non-linearities. The energy function is defined as $|\omega - 261.6Hz|$.*

any extra constraints inserted into the optimization formulation. Finally, we show images and quantitative error (see supplemental document for measurement details). We urge the reader to listen to the accompanying video for qualitative assessment. Material parameters for each example are listed in the accompanying supplemental material.

9.1 The Zoolophone

The Zoolophone is a glockenspiel with the ordinary rectangular keys replaced with animal shapes.

Input and Output. For each Zoolophone key we specify a fundamental frequency (the note of the key) as well as several overtones. Our optimization seeks to match these notes while reducing the amplitudes of all other frequencies.

Parametrization for Frequency Optimization. Each Zoolophone piece is modeled using b-spline curves [Piegl and Tiller 1997], the control points of which become our $\boldsymbol{p}$ (Fig. 9a-b). We generate

Method	Iterations	Func. evals	Func. cost
Simulated Annealing	250	250	0.034
CMA-ES	34	206	0.040
LCS (finite diff.)	7	40	0.006

Table 1: *The number of iterations, function evaluations, and the function cost of SA, CMA-ES, and LCS on a 2D toric acoustic shape design problem (Fig. 8), minimizing $\left|\omega_1^* - 261.6Hz\right|$.*

3D geometry by extruding these 2D curves in z. If the number of control points in the initial shape is expressive, then we use a free-form-deformation grid to create a reduced space of shape parameters to make our optimization more efficient (Fig. 9c-d).

Parametrization for Amplitude Optimization. We perforate the object to control the amplitude of the frequency spectrum (§6). This is accomplished using standard CSG operations (Fig. 9(e)), with each perforation parameterized by location and radius. As mentioned previously, perforating an object allows us to simplify the amplitude modulation cost function. It also avoids excessive deformation of the geometry for the sake of amplitude control.

Fabricated Result. We create one key for each note in the C Major scale, namely the notes C, D, E, F, G, A, B, and C. Figure 9 shows the final shapes for each key, along with the user-defined frequency goals, the initial frequency spectra, the spectra computed using only linear scaling, the result of our method, and the spectra from a professionally-manufactured glockenspiel. To aid visual comparison, we overlay raw frequency plots with bar charts that highlight the seven largest peaks in the frequency spectra. Furthermore, Table 2 shows all frequency errors for these examples. In terms of fundamental frequency error, our algorithm outperforms isotropic scaling in all but one case for which we are within 0.2Hz. For overtones, we always outperform isotropic scaling, reducing error by $\approx 4\times$ to $\approx 70\times$ depending on the example. The main source of our error is due to fabrication: the waterjet used to manufacture the Zoolophone could chip pieces and had difficulty following non-smooth paths in the optimized geometry. A better understanding of the constraints of this fabrication method would improve our ability to constrain our optimization, thus improving our results.
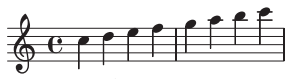
We also created several special zoolophone keys. The first two are elephants which demonstrated the ability of our algorithm to trade accuracy for shape preservation (Fig. 1: good shape and satisfactory acoustics; Fig. 9: more deformed shape ('anteater') but with better acoustics). See our supplemental video for a comparison. The next special key is a larger giraffe optimized to have a fundamental frequency at C4. Initial, isotropic scaling of the giraffe allowed us to match the fundamental frequency, but left a dense frequency spectrum (Fig. 12, bottom row). Our optimization scheme was able to significantly suppress unwanted frequencies, leading to a cleaner sound (Fig. 12, top row). Finally we demonstrate the ability of our method to control overtones, with a giraffe key that produces a chord when struck. This key has frequency peaks of equal amplitude at C, E, and G. We believe this is the first multi-tonal glockenspiel key ever produced (Fig. 4).

9.2 Tea for Three

Input and Output. Our method is not limited to producing 2D geometries. We perform frequency and amplitude modulation for bell-like cups, each optimized to produce a specific note, one of C, F or A. The cups were 3D printed by an online service.

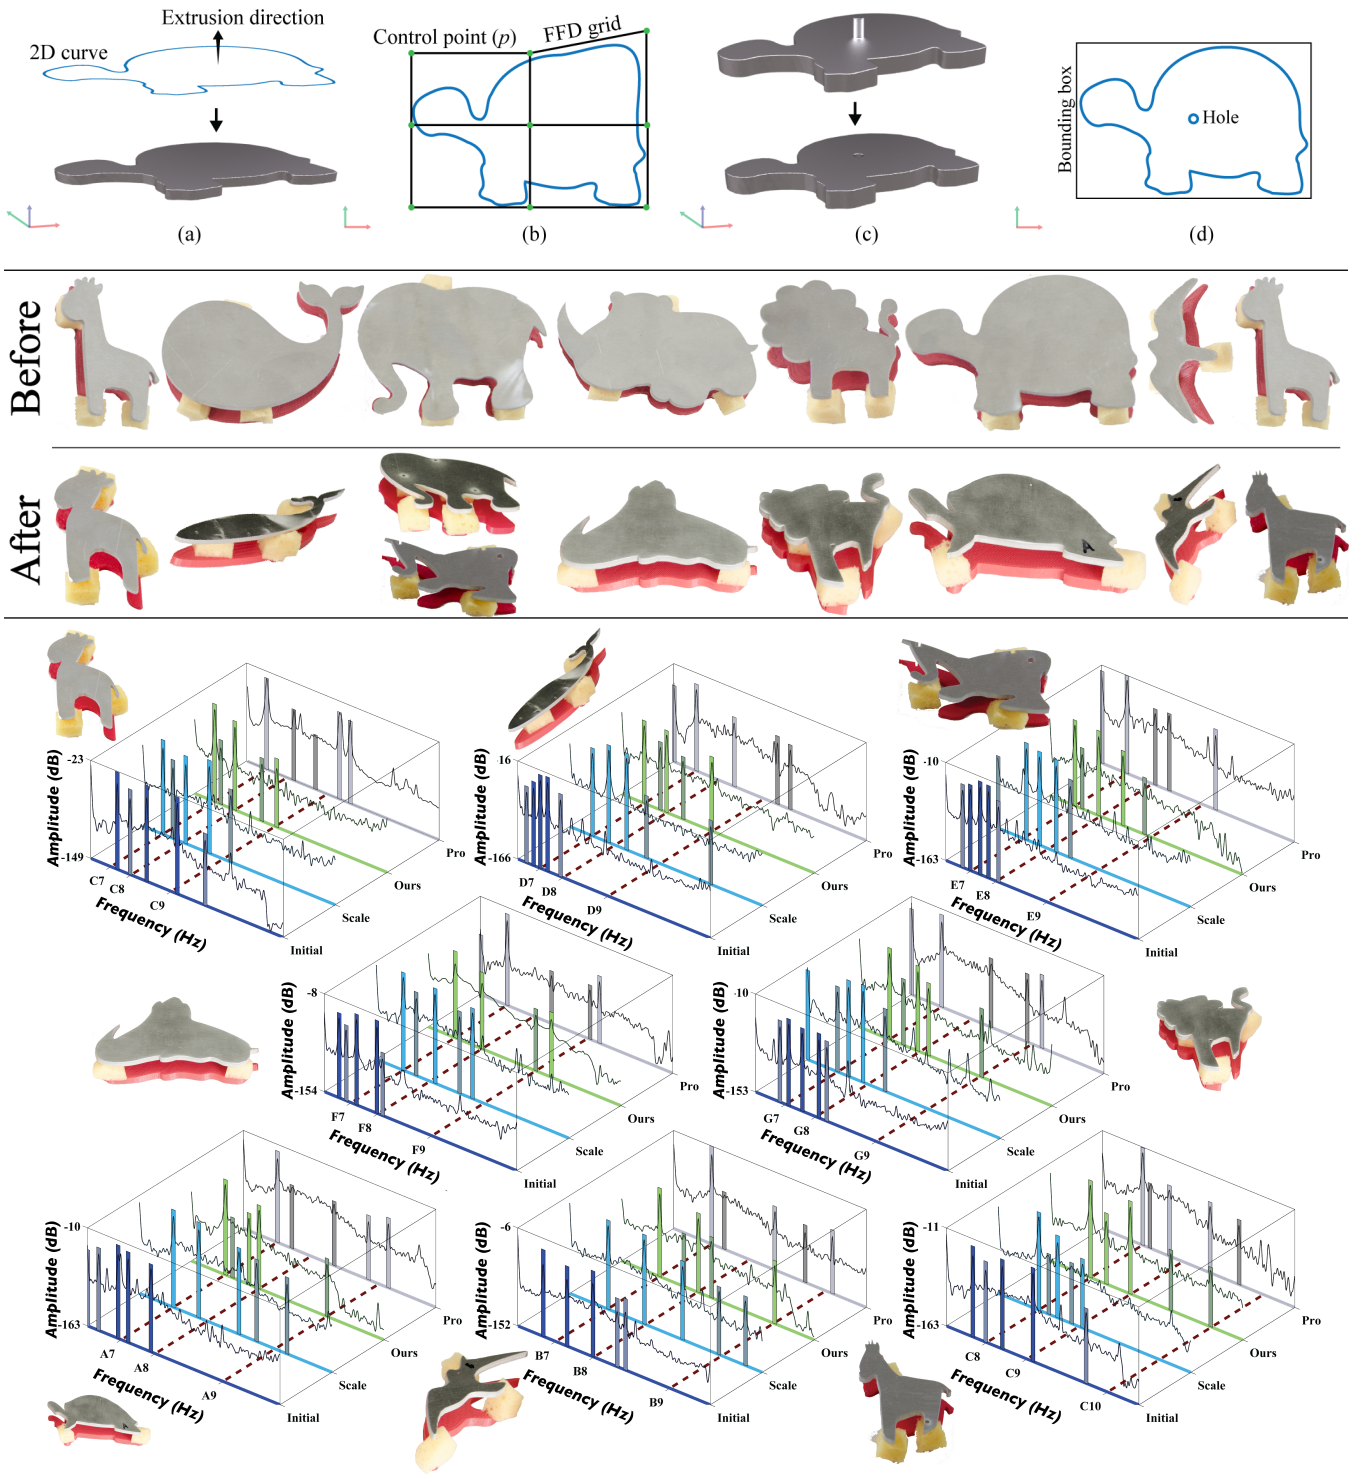

Figure 9: The Zoolophone. ***Top:*** *Frequency and amplitude parametrization: (a) Extrusion operation. (b) FFD parameterization (control points) and deformation. (c) Hole difference operation. (d) Amplitude-modulation parameterization.* ***Middle:*** *Shapes before and after optimization. The two 'after' elephants demonstrate a trade-off: (top) good shape with satisfactory acoustics, or (bottom) more deformed shape ('anteater') but with better acoustics.* ***Bottom:*** *Frequency response for each key in the zoolophone. We show four spectra for each key. The spectrum of the initial shape (dark blue), the spectrum after isotropic scaling to match the fundamental (light blue), the result of our method (green), and the spectrum of a professionally-manufactured glockenspiel (grey). The five largest frequency peaks for each spectrum are denoted by vertical bars. Goal notes are shown as dashed red lines along the ground plane and denoted using scientific pitch notation. Note that our goal is not to match the overtones of the professional key, only the fundamental frequency, and so please consider any match/non-match here to be a coincidence. The reader is encouraged to zoom in to examine each plot.*

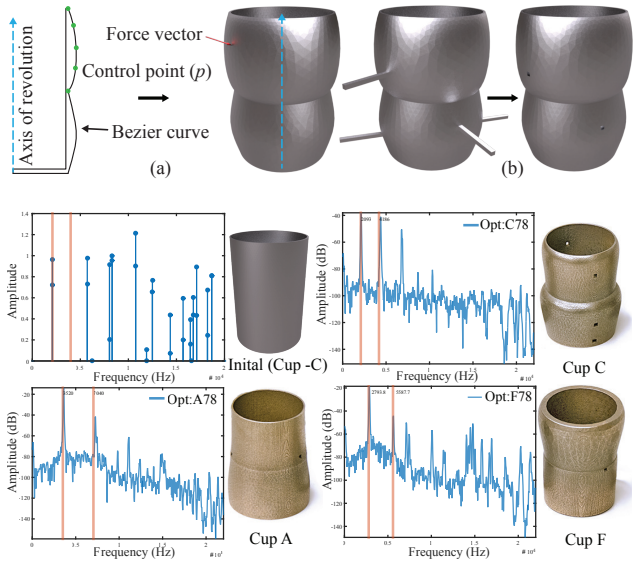

Figure 10: *Tea cups example.* ***Top****: Shape parameterization as a surface of revolution (a). Amplitude-modulation perforation (b).* ***Bottom****: Frequency spectra for the initial cup shape (simulated) as well as the three manufactured cups. The two goal frequencies are shown in the top right of each plot in scientific pitch notation.*

Parametrization for Frequency Optimization. We use surfaces of revolution to create our cup shapes (Fig. 10, first row). Here, the shape parameters for frequency optimization are the control points of cubic b-spline curves. These control points determine the shape change of the cup, as shown by green dots in Figure 10. Augmenting the parameterization with a wall thickness parameters gives us printable 3D geometry.

Parametrization for Amplitude Optimization. As described in §6, we first decide where generally we would like to hit the object to produce the desired frequency. Figure 10(a) shows vertices where the force is applied. Similar to what is described in §9.1, in 3D we achieve amplitude modulation by adding holes to the cup geometry. To this end, we perform a difference operation between the cup geometry and thin pipe-like cuboids (Figs. 10(b) & 10(c)). Alternatively, we could add holes using cylinders; however, difference operation with cylinders leads to unnecessary high tetrahedral count, with no added advantage for modulation. Finally, the parameters for energy described by Equation (5) are the width, thickness, angle-of-rotation, and height of the cuboids. All parameters had box constraints.

Fabricated Result. Using the same initial parameter values, we optimize for three different frequencies: C 7, 8; F 7, 8; A 7, 8 (Fig. 1). As with all 2D pieces, we optimize for and fabricate stands for the cups. However, our C cup has no active vibration modes on its bottom surface and so does not require a stand. Of all our results, the cups have the most fabrication errors. Two of the returned cups had significantly different densities than what was specified by the manufacturer. Despite this, the final products have average error $\approx 1.5\%$.

9.3 Cityscape

Input and Output. Thus far, we have optimized 2D and 3D object frequency responses for a single contact location. Next, we attempt the more challenging problem of a piece which produces two different notes when struck in two different locations.

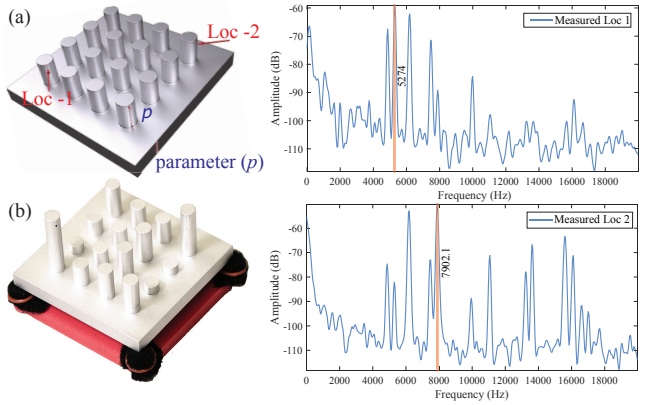

Figure 11: Cityscape. *(a) Initial design space and its parameterization. The location-arrows represent the contact-force locations and directions. (b) Fabricated result.*

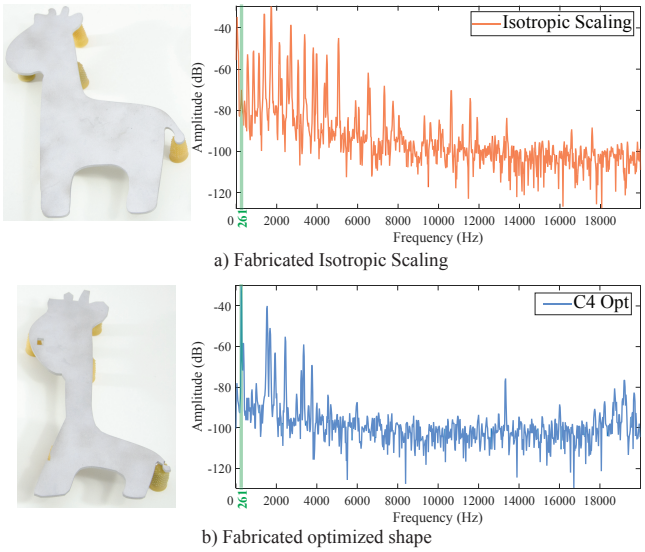

Figure 12: *We optimize a larger giraffe with a lower fundamental frequency (C4). Top: Isotropic scaling yields a dense spectrum of higher frequencies. Here, the amplitude of the fundamental frequency is diminished. Bottom: Our algorithm not only matches the fundamental frequency correctly at 261 Hz, but suppresses the amplitude of the remainder of the frequencies. The reader is encouraged to zoom in to examine both plots.*

Parametrization for Frequency Optimization. Our design space is a grid of vertical cylinders. The design parameters, $\boldsymbol{p}$, are the heights of the sixteen cylinders and the thickness of the base (Fig. 11(a)). The vertical distance between the top of the cylinder and the contact point remains constant.

Optimization Formulation. We must modify our standard cost functions (Eqs. (4) and (5)) to solve this multi-impact problem. Given the two desired frequencies in the spectrum and their respective impact locations (l_1, l_2), we solve the modified frequency optimization problem:

$$E(\boldsymbol{p}) = E_\omega(\boldsymbol{p}) + \lambda_1 \alpha_2(I_1) + \lambda_2 \alpha_1(I_2), \qquad (9)$$

where $\alpha_i(I_j)$ $i, j \in \{1, 2\}$ is the amplitude of the i^{th} user defined frequency at the j^{th} impact location, and λ_1 and λ_2 are user-defined weights. Our aim to have only one frequency be heard at

Shape	Note	Fundamental				Overtone 1			Overtone 2		
		Initial	Iso-scale	Full	Professional	Initial	Iso-scale	Full	Initial	Iso-scale	Full
Giraffe	C 7, 8, 9	32.20	0.81	0.04	0.51	2.00	7.38	0.31	7.70	9.50	1.85
Whale	D 7, 8, 9	30.78	0.37	0.56	0.37	5.82	13.61	0.54	3.30	5.86	0.32
Elephant	E 7, 8, 9	32.84	0.87	0.26	0.22	15.37	17.63	3.29	8.80	1.56	0.60
Rhino	F 7, 8, 9	45.27	0.18	0.18	0.18	1.92	9.25	0.00	4.53	0.96	0.00
Lion	G 7, 8, 9	18.46	0.35	0	0.44	5.10	7.68	2.70	4.73	17.54	1.19
Tortoise	A 7, 8, 9	66.96	0.62	0.02	0.02	6.25	12.52	0.17	7.66	13.43	0.02
Bird	B 7, 8, 9	32.82	1.08	0.40	0.68	0.87	14.16	1.02	1.79	4.43	0.71
Giraffe	C 8, 9, 10	33.77	0.57	0.21	0.57	7.76	26.24	0.65	13.16	18.43	0.49

Table 2: *Frequency percent errors for fabricated zoolophone pieces: for the initial shape, our preliminary isotropic scaling to match just the fundamental frequency, and our full non-linear optimization. For the fundamental frequency, we include a comparison to the professional glockenspiel; however, beyond this, the overtones are not intended to match.*

Shape Name	Notes	Freq. dims	Ampl. dims
Cup	F 7, 8	11	12
Cup	A 7, 8	11	12
Cup	C 7, 8	11	12
Cityscape	E 8 & B 8	17	-
Giraffe	C 7, 8, 9	18	12
Whale	D 7, 8, 9	18	12
Elephant	E 7, 8, 9	18	12
Rhino	F 7, 8, 9	18	12
Lion	G 7, 8, 9	18	12
Tortoise	A 7, 8, 9	18	12
Bird	B 7, 8, 9	18	12
Giraffe	C 8, 9, 10	18	12
Giraffe	C 4	18	12
Giraffe	C 6, E 6, G 6	18	12
Giraffe	C 7, E 7, G 7	18	12

Table 3: *Dimensionality of the optimized parameter space for frequency and amplitude stages for each shape.*

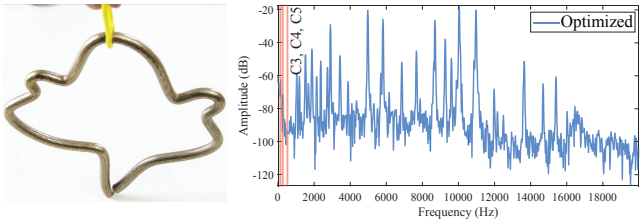

(a) Triangle Chime (b) Sound Spectrum

Figure 13: *We attempt to optimize a standard triangle, which typically has a very dense spectrograph, to produce note C3 with overtones of C4 and C5. While some peaks are matched, even drastic shape change (left) cannot dampen all other peaks in frequency.*

a time is equivalent to ensuring that the mode shape associated with each frequency has no displacement at the alternate impact point. It is this relationship that Equation 9 encodes.

Fabricated Result. Figure 11(b) shows the optimized fabricated result. We optimized hit-location 1 ($l1$) for ♪E8 (5274 Hz) and hit-location 2 ($l2$) for ♪B8 (7902.1 Hz). The recorded average percentage error was 0.2% for the two frequencies.

10 Conclusion

We have presented a functional specification algorithm for acoustic shape design. We have applied our new method to produce a variety of 2D and 3D metallic objects with user-defined impact frequency spectra. Using our techniques, new aesthetic-acoustic design is available for musicians, artists, and engineers to explore. If we consider that in the future our approaches could be applied to problems such as creating dampened operating mechanisms, then our work holds promise as a tool to help solve a broader class of sound control problems.

Limitations and Future Work. The main limitation of our technique is that the achievable control for objects with dense frequency spectra is limited. The C4 giraffe (Fig. 12) shows that our algorithm can perform simple frequency and amplitude control under these conditions. However, it was impossible to optimize for more than one frequency given the dimensionality of our shape parameterization. A similar failure case is shown in Figure 13, where we are unable to achieve control over the dense spectrum of a triangle chime. As a rule of thumb, denser spectrum control appears to come at the expense of geometry control, leading to undesirable results when our parameterizations are to restrictive. This restriction extends to overtone control: we were able to produce single key chords for C6E6G6 and C7E7G7, but not for C5E5G5 due to the increased spectrum density.

In some cases such as our elephant example, our algorithm could produce an acceptable acoustic result at the expense of unacceptable object geometry distortion. In these cases, adding additional user-defined shape constraints could help, though it is unclear how this would effect the fitted spectrum.

In general, choosing an appropriate parameterization is crucial for algorithm success since the parameterized geometry must be able to replicate the user-supplied frequency spectrum. Smart parameterization refinement algorithms, as well as fast methods for detecting "impossible" goals, are important open research questions.

Finally, our output is at the mercy of fabrication accuracy. Manufacturing our tea cups required multiple attempts due to inconsistencies in the 3D printing process which led to inconsistent material properties. Accounting for such variance during optimization is an important next step for fabrication algorithms. Further, while our method is rooted in the mechanical engineering approach of controlling specific object vibration modes, more perceptual metrics might allow us to avoid the mode matching cost functions and simplify the implementation.

Acknowledgements

We thank the anonymous reviewers for their feedback. A warm thank you to Sebastian Joseph, Emma Steinhardt, Henry Kavle, and Javier Ramos for fabricating the shapes, without which this paper would not be possible. Thank you to Joseph Booth for

the robot work. This research was supported in part by the National Science Foundation: CAREER-1453101, IIS-1116619, IIS 1447344; by the Air Force Research Laboratory and DARPA Memex program, and by generous gifts from Intel. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of funding agencies or others.

References

ANGELL, T., JIANG, X., AND KLEINMAN, R. 1997. A distributed source method for inverse acoustic scattering. *Inverse Problems 13*, 2, 531.

ARTHUR, D., AND VASSILVITSKII, S. 2007. k-means++: The advantages of careful seeding. In *Proceedings of the 18th annual ACM-SIAM symposium on Discrete algorithms*, 1027–1035.

BÄNGTSSON, E., NORELAND, D., AND BERGGREN, M. 2003. Shape optimization of an acoustic horn. *Computer methods in applied mechanics and engineering 192*, 11, 1533–1571.

BARBIERI, R., AND BARBIERI, N. 2006. Finite element acoustic simulation based shape optimization of a muffler. *Applied Acoustics 67*, 4, 346–357.

BARDENET, R., AND KÉGL, B. 2010. Surrogating the surrogate: accelerating gaussian-process-based global optimization with a mixture cross-entropy algorithm. In *27th International Conference on Machine Learning (ICML 2010)*, Omnipress, 55–62.

BICKEL, B., BÄCHER, M., OTADUY, M. A., MATUSIK, W., PFISTER, H., AND GROSS, M. 2009. Capture and modeling of non-linear heterogeneous soft tissue. *ACM Trans. Graph. 28*, 3 (July).

BICKEL, B., BÄCHER, M., OTADUY, M. A., LEE, H. R., PFISTER, H., GROSS, M., AND MATUSIK, W. 2010. Design and fabrication of materials with desired deformation behavior. *ACM Trans. Graph. 29*, 4 (July), 63:1–63:10.

BICKEL, B., KAUFMANN, P., SKOURAS, M., THOMASZEWSKI, B., BRADLEY, D., BEELER, T., JACKSON, P., MARSCHNER, S., MATUSIK, W., AND GROSS, M. 2012. Physical face cloning. *ACM Trans. on Graphics (SIGGRAPH 2012) 31*, 4 (July), 118:1–118:10.

BISHOP, C. M., ET AL. 2006. *Pattern recognition and machine learning*, vol. 1. springer New York.

BRANKE, J., DEB, K., MIETTINEN, K., AND SLOWINSKI, R. 2008. *Multiobjective Optimization: Interactive and Evolutionary Approaches*. LNCS sublibrary: Theoretical computer science and general issues. Springer.

CHADWICK, J. N., AN, S. S., AND JAMES, D. L. 2009. Harmonic shells: a practical nonlinear sound model for near-rigid thin shells. *ACM Transactions on Graphics (TOG) 28*, 5, 119.

CHAIGNE, A., AND DOUTAUT, V. 1997. Numerical simulations of xylophones. i. time-domain modeling of the vibrating bars. *The Journal of the Acoustical Society of America 101*, 1, 539–557.

CHEN, D., LEVIN, D. I. W., DIDYK, P., SITTHI-AMORN, P., AND MATUSIK, W. 2013. Spec2fab: A reducer-tuner model for translating specifications to 3d prints. *ACM Trans. Graph. 32*, 4.

CHOI, K. K., AND KIM, N.-H. 2006. *Structural sensitivity analysis and optimization 1: linear systems*. Springer Science & Business Media.

COOK, R. D., ET AL. 2007. *Concepts and applications of finite element analysis*. Wiley.

DE POLI, G., PICCIALLI, A., AND ROADS, C., Eds. 1991. *Representations of Musical Signals*. MIT Press, Cambridge, MA, USA.

DEMPSTER, A. P., LAIRD, N. M., AND RUBIN, D. B. 1977. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society, Series B 39*, 1, 1–38.

DIEGEL, O., 2013. 3d printed alto saxophone.

DOKMANIĆ, I., PARHIZKAR, R., WALTHER, A., LU, Y. M., AND VETTERLI, M. 2013. Acoustic echoes reveal room shape. *Proceedings of the National Academy of Sciences 110*, 30, 12186–12191.

DONG, Y., WANG, J., PELLACINI, F., TONG, X., AND GUO, B. 2010. Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph. 29*, 4 (July), 62:1–62:10.

DÜHRING, M. B., JENSEN, J. S., AND SIGMUND, O. 2008. Acoustic design by topology optimization. *Journal of sound and vibration 317*, 3, 557–575.

ESSL, G., AND COOK, P. R. 1999. Banded waveguides: Towards physical modeling of bowed bar percussion instruments. In *Proceedings of the International Computer Music Conference (ICMC)*, 321–324.

FEIJÓO, G. R., OBERAI, A. A., AND PINSKY, P. M. 2004. An application of shape optimization in the solution of inverse acoustic scattering problems. *Inverse problems 20*, 1, 199.

FONTANA, F., AND ROCCHESSO, D. 1998. Physical modeling of membranes for percussion instruments. *Acta Acustica united with Acustica 84*, 3, 529–542.

GORDON, C., WEBB, D., AND WOLPERT, S. 1992. Isospectral plane domains and surfaces via riemannian orbifolds. *Inventiones mathematicae 110*, 1, 1–22.

GORDON, C., WEBB, D. L., AND WOLPERT, S. 1992. One cannot hear the shape of a drum. *Bulletin of the American Mathematical Society 27*, 1, 134–138.

HAFNER, C., MUSIALSKI, P., AUZINGER, T., WIMMER, M., AND KOBBELT, L. 2015. Optimization of natural frequencies for fabrication-aware shape modeling. In *ACM SIGGRAPH Posters*, ACM, New York, NY, USA, 82:1–82:1.

HANSEN, N., MULLER, S., AND KOUMOUTSAKOS, P. 2003. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation 11*, 1, 1–18.

HAŠAN, M., FUCHS, M., MATUSIK, W., PFISTER, H., AND RUSINKIEWICZ, S. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM Trans. on Graphics (SIGGRAPH 2010) 29*, 4 (July), 61:1–61:10.

ISHIGURO, Y., AND POUPYREV, I. 2014. 3d printed interactive speakers. ACM, CHI 2014, 1733–1742.

JAMES, D. L., BARBIČ, J., AND PAI, D. K. 2006. Precomputed acoustic transfer: output-sensitive, accurate sound generation for geometrically complex vibration sources. *ACM Transactions on Graphics (TOG) 25*, 3, 987–995.

KAC, M. 1966. Can one hear the shape of a drum? *American Mathematical Monthly*, 1–23.

KIRKPATRICK, S., VECCHI, M. P., ET AL. 1983. Optimization by simmulated annealing. *science 220*, 4598, 671–680.

LABELLE, F., AND SHEWCHUK, J. R. 2007. Isosurface stuffing: Fast tetrahedral meshes with good dihedral angles. *ACM Trans. Graph. 26*, 3 (July).

LEE, H.-P., AND LIN, M. C. 2012. Fast optimization-based elasticity parameter estimation using reduced models. *The Visual Computer 28*, 6-8, 553–562.

LI, S., HUANG, J., DE GOES, F., JIN, X., BAO, H., AND DESBRUN, M. 2014. Space-time editing of elastic motion through material optimization and reduction. *ACM Trans. Graph. 33*, 4 (July).

LLOYD, D. B., RAGHUVANSHI, N., AND GOVINDARAJU, N. K. 2011. Sound synthesis for impact sounds in video games. In *Symposium on Interactive 3D Graphics and Games*, ACM, PAGE–7.

MARBURG, S. 2002. Developments in structural-acoustic optimization for passive noise control. *Archives of computational methods in engineering 9*, 4, 291–370.

MATUSIK, W., AJDIN, B., GU, J., LAWRENCE, J., LENSCH, H. P., PELLACINI, F., AND RUSINKIEWICZ, S. 2009. Printing spatially-varying reflectance. *ACM Trans. Graphics 28*, 5 (Dec.).

MCKAY, M., BECKMAN, R., AND CONOVER, W. 2000. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics 42*, 1.

NAGARAJ, K. S. 2014. Stochastically constrained simulation optimization on mixed-integer spaces.

O'BRIEN, J. F., SHEN, C., AND GATCHALIAN, C. M. 2002. Synthesizing sounds from rigid-body simulations. SCA '2002, 175–181.

PAPAS, M., JAROSZ, W., JAKOB, W., RUSINKIEWICZ, S., MATUSIK, W., AND WEYRICH, T. 2011. Goal-based caustics. *Computer Graphics Forum 30*, 2, 503–511.

PETTERSSON, T. 2008. Global optimization methods for estimation of descriptive models.

PIEGL, L., AND TILLER, W. 1997. *The NURBS Book (2Nd Ed.)*. Springer-Verlag New York, Inc., New York, NY, USA.

RAGHUVANSHI, N., AND LIN, M. C. 2006. Interactive sound synthesis for large scale environments. In *Proceedings of the 2006 symposium on Interactive 3D graphics and games*, 101–108.

REN, Z., YEH, H., AND LIN, M. C. 2013. Example-guided physically based modal sound synthesis. *ACM Trans. Graph. 32*, 1 (Feb.).

SHABANA, A. A. 1995. *Theory of Vibration: An Introduction*. Mechanical Engineering Series. Springer.

SKOURAS, M., THOMASZEWSKI, B., BICKEL, B., AND GROSS, M. 2012. Computational design of rubber balloons. *Computer Graphics Forum 31*, 2pt4, 835–844.

SKOURAS, M., THOMASZEWSKI, B., COROS, S., BICKEL, B., AND GROSS, M. 2013. Computational design of actuated deformable characters. *ACM Trans. Graph. 32*, 4 (July), 82:1–82:10.

SNOEK, J., LAROCHELLE, H., AND ADAMS, R. P. 2012. Practical bayesian optimization of machine learning algorithms. In *Neural Information Processing Systems*.

UMETANI, N., MITANI, J., AND IGARASHI, T. 2010. Designing custom-made metallophone with concurrent eigenanalysis. In *Proceedings of the Conference on New Interfaces for Musical Expression (NIME)*, 26–30.

VAN DEN DOEL, K., AND PAI, D. K. 1998. The sounds of physical shapes. *Presence: Teleoperators and Virtual Environments 7*, 4.

WADBRO, E., AND BERGGREN, M. 2006. Topology optimization of an acoustic horn. *Computer methods in applied mechanics and engineering 196*, 1, 420–436.

WAMPLER, K., AND POPOVIĆ, Z. 2009. Optimal gait and form for animal locomotion. *ACM Trans. Graph. 28*, 3 (July), 60:1–60:8.

WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM Trans. on Graphics (SIGGRAPH 2009) 28*, 3 (July), 32:1–32:6.

WRIGHT, S., AND NOCEDAL, J. 1999. *Numerical optimization*, vol. 2. Springer New York.

XU, H., LI, Y., CHEN, Y., AND BARBIČ, J. 2015. Interactive material design using model reduction. *ACM Trans. Graph. 34*, 2 (Mar.).

YAMASAKI, S., NISHIWAKI, S., YAMADA, T., IZUI, K., AND YOSHIMURA, M. 2010. A structural optimization method based on the level set method using a new geometry-based re-initialization scheme. *International journal for numerical methods in engineering 83*, 12, 1580–1624.

YOO, H. H., CHO, J. E., AND CHUNG, J. 2006. Modal analysis and shape optimization of rotating cantilever beams. *Journal of Sound and vibration 290*, 1, 223–241.

YU, Y., JANG, I. G., KIM, I. K., AND KWAK, B. M. 2010. Nodal line optimization and its application to violin top plate design. *Journal of Sound and Vibration 329*, 22, 4785–4796.

YU, Y., JANG, I. G., AND KWAK, B. M. 2013. Topology optimization for a frequency response and its application to a violin bridge. *Structural and Multidisciplinary Optimization 48*, 3, 627–636.

ZELDITCH, S. 2000. Spectral determination of analytic bi-axisymmetric plane domains. *Geometric & Functional Analysis GAFA 10*, 3, 628–677.

ZHENG, C., AND JAMES, D. L. 2010. Rigid-body fracture sound with precomputed soundbanks. *ACM Transactions on Graphics (TOG) 29*, 4, 69.

ZHENG, C., AND JAMES, D. L. 2011. Toward high-quality modal contact sound. *ACM Transactions on Graphics (TOG) 30*, 4, 38.

ZHU, L., XU, W., SNYDER, J., LIU, Y., WANG, G., AND GUO, B. 2012. Motion-guided mechanical toy modeling. *ACM Trans. Graph. 31*, 6 (Nov.), 127:1–127:10.