

Point-based Acoustic Scattering for Interactive Sound Propagation via Surface Encoding

Hsien-Yu Meng, Zhenyu Tang, Dinesh Manocha

University of Maryland, College Park

{mengxy19,zhy,dmanocha}@umd.edu

Abstract

We present a novel geometric deep learning method to compute the acoustic scattering properties of geometric objects. Our learning algorithm uses a point cloud representation of objects to compute the scattering properties and integrates them with ray tracing for interactive sound propagation in dynamic scenes. We use discrete Laplacian-based surface encoders and approximate the neighborhood of each point using a shared multi-layer perceptron. We show that our formulation is permutation invariant and present a neural network that computes the scattering function using spherical harmonics. Our approach can handle objects with arbitrary topologies and deforming models, and takes less than 1ms per object on a commodity GPU. We have analyzed the accuracy and perform validation on thousands of unseen 3D objects and highlight the benefits over other point-based geometric deep learning methods. To the best of our knowledge, this is the first real-time learning algorithm that can approximate the acoustic scattering properties of arbitrary objects with high accuracy.

1 Introduction

Many interactive applications like games and virtual environments need to generate realistic or plausible sound effects. The resulting audio-visual renderings can enhance the realism and improve a user’s level of immersion [Liu and Manocha, 2020]. Current techniques used in game engines or VR systems for interactive sound rendering are based on reverberation filters or ray tracing. However, these methods are unable to generate realistic low-frequency sound effects corresponding to diffraction or interference.

The most accurate methods for sound propagation and rendering directly solve the acoustic wave-equation using numeric solvers. These techniques either compute the global acoustic pressure field for the entire scene [Raghuvanshi and Snyder, 2018] or compute the *acoustic scattering effects* of each separated object [Rungta *et al.*, 2018; Kuttruff, 2016]. The global methods are very expensive, have a high memory overhead, and are limited to static scenes.

Acoustic scattering corresponds to the disturbance of a given incident sound field due to an object’s shape and surface properties. It can be regarded as one of the fundamental characteristics of an object. The effect of scattering can be expressed in terms of a scattered sound field. The computation of acoustic scattering fields is also useful for other applications, including sound synthesis, sound source localization, object recognition, etc. The most accurate methods for acoustic scattering field computation use solvers like boundary-element methods. In practice, accurate computation of acoustic scattering field of each object can take a few minutes and cannot be directly used within an interactive application. Instead, some hybrid sound rendering systems precompute the scattering fields of rigid objects and combine them with interactive ray tracing at runtime. Recently, some machine learning based methods have been proposed to compute the scattering fields [Pulkki and Svensson, 2019; Fan *et al.*, 2020; Tang *et al.*, 2020]. However, they are limited to simple shapes (e.g., convex shapes) or only work well for low frequencies.

Main Results: We present novel techniques based on geometric deep learning on differential coordinates to approximate the acoustic scattering properties of arbitrary objects. Our approach is general and makes no assumption about object’s shape, genus, or rigidity. We approximate the objects using point clouds, and each point in the point cloud representation is encoded in a high-dimensional latent space. Moreover, the local surface shapes in the latent space are encoded using surface encoders. This enables us to handle arbitrary topology. Our network takes the point cloud as an input and outputs the spherical harmonic coefficients that represent the acoustic scattering field.

We present techniques to generate the ground truth data using an accurate wave-solver on a large geometric dataset. We have evaluated the performance on thousands of objects that are very different from the training database (with varying convexity, genus, shape, size, orientation) and observe high accuracy. We also perform an ablation study to highlight the benefits of our approach. The additional runtime overhead of estimating the scattering field from neural networks is less than 1ms per object on a NVIDIA GeForce RTX 2080 Ti GPU. We also prove that our learning method is permutation and rotation invariant, which is an important characteristic for accurate computation of acoustic scattering

fields. We have also integrated our geometric learning algorithm with ray tracing for hybrid sound propagation in interactive applications.

2 Related Works

2.1 Wave-Acoustic Simulation

Wave-acoustic simulation methods aim to solve the wave equation that governs the propagation and scattering of sound waves. Some conventionally used numeric methods include the finite-element method, the boundary-element method and the finite-difference time domain, as surveyed in [Liu and Manocha, 2020]. The common requirement of these methods is the proper discretization of the problem domain (e.g., spatial resolution, time resolution), which means the time complexity of them will scale drastically with the simulation frequency. An alternative way to use wave simulation results in real-time applications is to pre-compute a large amount of sound fields in a scene and compressing the results using perceptual metrics [Raghuvanshi and Snyder, 2018]. However, they are limited to static scenes.

2.2 Learning-Based Acoustics

There is considerable work on developing machine learning methods for applications corresponding to audio-visual analysis and acoustic scene classification [Bianco *et al.*, 2019]. In comparison, there are much fewer works studying the generation of acoustic data from a physical perspective. [Pulkki and Svensson, 2019] propose to use a neural network to model the acoustic scattering effect from rectangular plates. Recently, [Fan *et al.*, 2020] train convolutional neural networks (CNNs) to learn to map planar sound fields induced by convex scatterers in 2D. [Tang *et al.*, 2020] uses the well-known PointNet method to compute the scattering properties for point clouds, but it only works well for low frequencies. Motivated by the last method, we aim to develop networks that can deal with object geometry in 3D and extend the generality of this learning-based approach.

2.3 Geometric Deep Learning and Shape Representation

There is considerable recent work on generating plausible shape representations for 3D data, including voxel-based [Zhou and Tuzel, 2018], point-based [Charles *et al.*, 2017; Qi *et al.*, 2017; Wang *et al.*, 2019] and mesh-based [Hanoeka *et al.*, 2019] shape representations. This includes work on shape representation by learning implicit surfaces on point clouds [Smirnov *et al.*, 2020], hierarchical graph convolution on meshes [Mo *et al.*, 2019], encoding signed distance functions for surface reconstruction [Park *et al.*, 2019], etc. However, previous methods on point cloud shape representations learn by designing loss functions to constrain surface smoothness on global Cartesian coordinates. Such functions only provide spatial information of each point and lack information about local shape of the surface compared to the explicit discretization of the continuous Laplace-Beltrami operator and curvilinear integral. Instead, we use point cloud based learning algorithms, which do not require mesh Laplacians

for graph neural networks. This makes our approach applicable to all kind of dynamic objects, including changing topologies.

3 Our Approach

Our goal is to directly predict the acoustic scattering field (ASF) resulting from an incident wave and a sound scatterer of known shape (described by a triangle mesh). A set of frequency dependent ASFs are visualized in Figure 1 on a disk plane. The two key steps are: (i) defining a compact and efficient way to encode the ASF; (ii) designing a neural network architecture that can effectively exploit the 3D geometric information so as to establish the mapping between the object shape and the ASF. In this section, we provide necessary background on wave acoustics and present an overview of our approach.

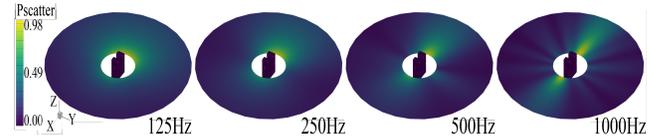


Figure 1: Acoustic scattering of the same object at four different frequencies, assuming a plane wave travelling to the $-x$ direction and a sound-hard boundary condition for the object/scatterer in the center. The scattering patterns are significantly different among frequency bands. Our learning method can compute the scattering fields with high accuracy for arbitrary objects at interactive rates, including at higher frequencies.

3.1 Acoustic Scattering Fields

We are interested in knowing how the incident sound field will be scattered when it acts on an object surface. Let the sound pressure at location \mathbf{x} and time t be $P(\mathbf{x}, t)$, and the wave equation describes the propagation of sound:

$$(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2})P(\mathbf{x}, t) = 0, \quad (1)$$

where c is the speed of sound, which is commonly assumed to be constant in non-large scale indoor environments. Transforming into the frequency domain, we obtain the homogeneous Helmholtz equation:

$$(\nabla^2 + k^2)p(\mathbf{x}, \omega) = 0, \quad (2)$$

where ω is the wave frequency and $k = \frac{\omega}{c}$ is the wavenumber. When a specific boundary condition is given, its solution in 3D spherical coordinates can be decomposed as

$$p(r, \theta, \phi, \omega) = \sum_{l=0}^{\infty} h_l(kr) \sum_{m=-l}^{+l} c_l^m(\omega) Y_l^m(\theta, \phi), \quad (3)$$

where $h_l(kr)$ is the Hankel function which encodes the radial pressure changes, $Y_l^m(\theta, \phi)$ is the spherical harmonics term which represents the angular part of the solution, and $c_l^m(\omega)$ is the spherical harmonics coefficients for frequency ω . For a maximum order L , there are $(L + 1)^2$ coefficients. Since the Hankel and the Spherical Harmonics functions are easy to evaluate, we only need to find the coefficients $c_l^m(\omega)$ to fully describe the acoustic scattering field.

3.2 Network Inference

Our goal is to use an appropriate geometric representation for the underlying objects in the scene so that we can apply geometric deep learning methods to infer the acoustic scattering field from the object shape. It is also important that our approach should be able to handle highly dynamic scenes with general objects (i.e., no assumption on their convexity or genus or topology). Moreover, our approach should be able to generate accurate ASFs results for changing topologies. The input to our system is the point cloud representation ($N \times 3$ matrix, N being the number of 3D points) of an object. We represent each point and its local surface by a higher dimensional surface in the latent space formed by a surface encoder. Based on Kolmogorov–Arnold representation theorem, we can represent the pressure field function as a function defined on a set of points ($N \times 3$), which is a multivariate continuous function composed of continuous functions of one variable encoded in a higher dimension. Moreover, borrowing the power of the *De Finetti Theorem*, our pressure field function is further formulated as Eq. 14. Please refer to § 4 and § A for more details of our design. The desired output is the spherical harmonic coefficients vector up to order $L = 3$ described in § 3.1. In practice, acoustic scattering corresponding to different frequencies can exhibit different distributions and the computational complexity increases as a cubic function of the frequency. Therefore, we train several networks for different frequency bands and in this paper we present the results for the frequency range $125\text{Hz} \sim 1000\text{Hz}$, which is sufficient for low-frequency sound propagation effects. In theory, our approach can extend to higher frequencies, though the cost of generating the training data would increase significantly. We present a novel network architecture to compute these characteristics for arbitrary or deforming objects.

4 Acoustic Scattering via Surface Encoding

In this section we introduce our shape Laplacian based point representation using surface encoders. For predicting ASF, the detailed geometric characteristics of object representation play an essential role with respect to the simulation frequency. This motivates our design in terms of using differential coordinates [Sorkine, 2006], which describe the discretization of the curvilinear integral around a given point (i.e., encodes the mean curvature). Moreover, we also provide the proof of its permutation invariant property via Newton’s identities in the appendix, as that is important in the context of our sound propagation algorithm. Finally, following De Finetti theorem [Heath and Sudderth, 1976], we justify our design of the surface encoder and show its benefits in terms of accurately computing the acoustic scattering functions for arbitrary objects.

4.1 Local Surface Shape

Previous works on point cloud learning algorithms mostly focus on designing per-point operations [Charles *et al.*, 2017], encoding per-point features to estimate continuous shape functions [Xu *et al.*, 2019], or minimizing loss between a point normal vector and its connected vertices [Liu *et*

al., 2019]. We extend these ideas to compute the acoustic scattering functions.

For each point in the input cloud and its neighborhood in the Euclidean space, we assume that it can form a piecewise smooth surface around the point. Each point is encoded by the shared multi-layer perceptron (MLP) [Rumelhart *et al.*, 1985] and can be represented by a vector in the higher dimensional latent space $\mathcal{Z} \in \mathcal{R}^{128}$ (see Figure 2) Thus, a piecewise-linear approximation of the surface around a given point can be used to estimate the local surface shape, where the differential coordinate [Sorkine, 2006] (i.e., δ -coordinates) of each vertex v_i can be expressed as:

$$\vec{\delta}_i = \frac{1}{d_i} \sum_{j \in \mathcal{N}(i)} (\vec{v}_i - \vec{v}_j). \quad (4)$$

Here δ_i encapsulates the local surface shape, $\mathcal{N}(i)$ represents the K nearest neighbors of vertex v_i in the Euclidean space, and $d_i = |\mathcal{N}(i)|$ is the number of immediate neighbors of v_i . To estimate the mean curvature of the local surfaces formed by each point and its spatial neighbors, we use the radial basis function (RBF) $\varphi(\cdot) = \exp^{-\|\cdot\|^2}$ to weigh each vector, rather than using the uniform weight shown in Equation (4). Since there are $N!$ permutations for a point cloud with N vertices, every operation on point clouds should be permutation invariant (i.e., the permutation of input points should not change the output of our network). Therefore, we design an algorithm based on discrete laplacian and prove its permutation invariance property in § 4.4.

Our weight function is designed to be positive definite and symmetric for any choice of data coordinates.

4.2 Discrete Laplacian Based Surface Encoder

To encapsulate the local surface shape, each point v_i is projected onto a higher dimensional space using MLPs, and the high-dimensional surface is defined on the latent space z as shown on the top of Figure 2.

Specifically, for one layer MLP, the representation is $z_i = \text{relu}(\vec{w} \cdot \vec{\delta}_i + b)$, where \vec{w} and b are learnable parameters in our network. To calculate the δ -coordinates of a point v_i and the closed simple surface curve around it, its immediate neighbors in the Euclidean space (i.e., $\mathcal{N}(v_i)$) is used to evaluate Equation (5):

$$\vec{\delta}_i^*(v_i) = \sum_{j \in \mathcal{N}(i)} \frac{\varphi(\vec{v}_i - \vec{v}_j)(\vec{v}_i - \vec{v}_j)}{\sum_{j \in \mathcal{N}(i)} \varphi(\vec{v}_i - \vec{v}_j)}. \quad (5)$$

In the latent space, the local surface shape is encoded as a high-dimensional surface where $\mathcal{N}(z_i)$ is evaluated. The direction of the differential coordinate vector, also defined in Equation (5), approximates the local normal direction.

In order to compare the two δ -coordinate representations, we highlight the error in decibel (dB) between the pressure fields reconstructed from groundtruth spherical harmonics term and the predicted ones using different neural networks in Table 1. We observe that δ^* -coordinates result in smaller errors. This indicates that our formulation provides a better approximation of ASFs.

4.3 Neural Network Design

Our neural network takes the point cloud as an $N \times 3$ input, where N represents the number of points in the point cloud. The output is the spherical harmonic coefficients c_l^m , $-l \leq m \leq l, 0 \leq l \leq 3$, resulting in a vector of length 16. The network is illustrated in Figure 2.

For each point v_i in the $N \times 3$ point cloud, the discrete laplacian is evaluated on its immediate neighbors $\mathcal{N}(v_i)$ according to Equation (5). Then one convolutional layer and three MLP layers are used to encode the piecewise-smooth local shape around v_i into a latent space $\mathcal{Z} \subset \mathcal{R}^{128}$. For each $z_i \in \mathcal{Z}$, the discrete laplacian Equation (5) is further evaluated where $\mathcal{N}(v_i)$ (i.e., neighbors of v_i in the Euclidean space \mathcal{R}^3) and $\mathcal{N}(z_i)$ (i.e., neighbors of z_i in the latent space \mathcal{Z}^{128}) scale the δ -coordinates \mathcal{R}^{128} , respectively. This results in a $2K \times 128$ matrix. The final representation of v_i , *feature*(v_i) is:

$$\begin{pmatrix} z_i^T, \frac{\varphi(\vec{v}_i - \vec{v}_1)(\vec{z}_i - \vec{z}_1)^T}{\sum_{j \in \mathcal{N}(v_i)} \varphi(\vec{v}_i - \vec{v}_j)}, \dots, \frac{\varphi(\vec{v}_i - \vec{v}_K)(\vec{z}_i - \vec{z}_K)^T}{\sum_{j \in \mathcal{N}(v_i)} \varphi(\vec{v}_i - \vec{v}_j)}, \\ \frac{\varphi(\vec{z}_i - \vec{z}_1)(\vec{z}_i - \vec{z}_1)^T}{\sum_{j \in \mathcal{N}(z_i)} \varphi(\vec{z}_i - \vec{z}_j)}, \dots, \frac{\varphi(\vec{z}_i - \vec{z}_K)(\vec{z}_i - \vec{z}_K)^T}{\sum_{j \in \mathcal{N}(z_i)} \varphi(\vec{z}_i - \vec{z}_j)} \end{pmatrix}, \quad (6)$$

where $K = 5$ is the number of the nearest neighbors and 128 is the dimension of the latent space. Now we have the piecewise approximation of the local shape around v_i as described in Equation (6), which is composed of the center point $z_i \in \mathcal{R}^{128}$ and the $2K \times 128$ weighted discrete laplacian matrix. The vectors are concatenated and convolved with kernels of size $[1, (1 + 2K)]$, followed by two MLP layers and four fully connected layers with *tanh* activation functions since the value of c_l^m ranges from $[-1, 1]$. The output of the network is the spherical harmonics coefficients c_l^m and the loss function is the \mathcal{L}_2 norm of the difference between the predicted c_l^m and the ground truth calculated using *FastBEM Acoustics solver*¹.

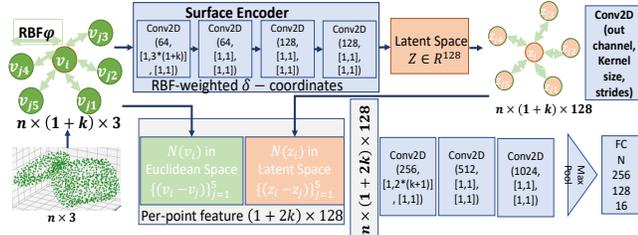


Figure 2: **Network architecture:** The input of our neural network is a $N \times 3$ point cloud ($N = 1024$) and the output of our network is the spherical harmonic coefficients as a vector of length 16 (up to 3rd order spherical harmonics). The piecewise smooth local shape around every point in the point cloud is encoded by virtue of the surface encoder, forming a high dimensional latent space in \mathcal{R}^{128} . The RBF-weighted δ -coordinates (Equation (5)) where neighbors are found in the Euclidean space and the latent space are concatenated to represent the curvature feature, forming Equation (6). We leverage the geometric details by several Conv2D and MLP layers and output 16 spherical harmonics coefficients c_l^m .

¹<https://www.fastbem.com/>

4.4 Discrete Laplacian and Permutation Invariance by De Finetti Theorem

In this section we describe the permutation invariant property of the RBF-weighted discrete Laplacian Equation (6). We mainly discuss the function composed of vectors in the high dimensional latent space projected by the surface encoder shown on the top of Figure 2.

Let $f_i(v_1, \dots, v_N) = \sum_{j \in \mathcal{N}(v_i)} \psi(v_i - v_j)$, where ψ is a continuous function composed of MLP layers. by the nature of *Kolmogorov–Arnold representation theorem* (or superposition theorem) and multi-layer neural networks [Kůrková, 1992], every multivariate continuous function can be represented as a superposition of continuous functions of one variable. Thus, $f_i(v_1, \dots, v_N)$ can be rewritten as: $\Phi(\sum_{j=1}^N \Psi(v_j))$, where $\Psi : \mathcal{R}^3 \rightarrow \mathcal{R}^{2N+1}$ and $\Phi : \mathcal{R}^{2N+1} \rightarrow \mathcal{R}^3$. In our work, we do not project v onto \mathcal{R}^{2N+1} ; instead, we encode v as $z \in \mathcal{R}^{128}$. First, we show that f_i can be regarded as a sum-of-power mapping defined by the coordinate functions. Second, we show that such mapping is *injective* and has continuous inverse mapping. The domain of f_i is a compact set, and thereby the image of f_i is a compact set as well. Thus, f_i is a homeomorphism between its domain and its image, and thereby is symmetric and homogeneous as per Newton’s identities (details in Appendix A.1). Further, we demonstrate that f_i is a *permutation invariant continuous* function since it can be considered as compositions of sum-of-power mappings.

Consider $\Psi(v) := [1, v, v^2, \dots, v^M], \Psi : \mathbb{R}^3 \rightarrow \mathbb{R}^{M+1}$, where M is the dimension of our surface encoder space. Let $E(v) = \sum_{j=1}^M \Psi(v_j)$, then

$$\begin{aligned} E(v) &= [M, \sum_{m=1}^M v_m, \sum_{m=1}^M v_m^2, \dots, \sum_{m=1}^M v_m^M] \\ &= [M, p_1(v_1, v_2, \dots, v_M), p_2(v_1, v_2, \dots, v_M), \\ &\dots, p_M(v_1, v_2, \dots, v_M)], \quad (7) \\ &= [E_0(v_1, v_2, \dots, v_M), E_1(v_1, v_2, \dots, v_M) \\ &\dots, E_M(v_1, v_2, \dots, v_M)], \end{aligned}$$

where p is the *power-sums* (defined in Equation (10), Appendix A.2). We observe that p is symmetric (i.e., does not change if we permute v_1, \dots, v_M) and homogeneous of degree M . E_q is defined in Lemma 4.1. According to Lemma 4.1, each element in $E(v)$ is injective. According to Lemma A.2 (Appendix A), $E(v)$ has inverse continuous mapping. Therefore, E is a homeomorphism between \mathcal{R}^3 and \mathcal{R}^{M+1} . Since Φ is the composition of several continuous functions ψ , this implies the continuity of Φ .

Lemma 4.1. Let $\mathbb{X} = \{(x_1, \dots, x_M) \in [0, 1]^M : x_1 \leq x_2 \leq \dots \leq x_M\}$. The *sum-of-power mapping* $\mathbb{E} : \mathbb{X} \rightarrow \mathbb{R}^{M+1}$ defined by the coordinate functions $:E_q(X) := \sum_{m=1}^M (x_m)^q, q = 0, \dots, M$ is *injective*.

The proof of Lemma 4.1 is in Appendix A. We can conclude that there exists functions Φ and Ψ such that f_i is permutation invariant as long as Ψ and Φ are continuous

functions. Our idea is to approximate the *scattering field* defined by a set of points (i.e., the closed simple surface curvature of a given point in the point cloud). Based on De Finetti theorem, we can derive the acoustic scattering function in Eq. 15 (Appendix A.2). We further highlight the benefit of our Discrete Laplacian based surface encoder by showing that it can be regarded as sum of power mappings. Thereby, Eq. 15 holds as per its continuity.

5 Evaluation and Application

In this section, we describe how our network (shown in Fig. 2) is used to approximate the ASF of an object and use that ASF for interactive sound propagation.

5.1 Data Generation and Training

We need a sufficiently large dataset for training and evaluating our method. We sample 100,000 3D objects from the *ABC Dataset* [Koch *et al.*, 2019]. All mesh files are pre-processed to be randomly rotated, scaled between $1 \sim 2m$ and centered. The *FastBEM Acoustics* is used to generate accurate acoustic scattering fields. Specifically, we simulate a plane wave travelling to the $-x$ axis direction, and all objects are assumed to have a zero Neumann boundary condition to solve Equation (2). While other boundary conditions are possible, we do not extensively study this variable in this work. After the acoustic scattering field has been solved at frequencies bands of $\{125, 250, 500, 1000\}Hz$, we use the *pyshools*² package to project the directional field to spherical harmonic coefficients c_l^m . By using a maximum order of $L = 3$, we are able to maintain a relative projection error below 2%. The original object meshes are also converted to point clouds using furthest sampling for 1024 points. The data generation pipeline takes about 12 days using a desktop with 32 Intel(R) Xeon(R) Gold 5218 CPU cores.

We split our dataset into training, validation and test set following a 8:1:1 ratio. Our neural network is trained on an NVIDIA GeForce RTX 2080 Ti GPU and takes less than 1ms per object for inference. The Adam optimizer is used and it decays exponentially with decay rate and decay step equal to 0.9 and $10 \times \#training$ examples, respectively.

5.2 ASF Computation: Results and Evaluation

We evaluate our network performance on thousands of unseen objects from the *ABC dataset*, perform ablation study on RBF-weighted δ -coordinates and implicit surface encoders, and show that our network is robust to Gaussian noise ($\sim N(0, 0.05)$). We compare our results with PointNet [Charles *et al.*, 2017] and DGCNN [Wang *et al.*, 2019] in Table 1 and Table 2, respectively. We justify the design of our network with this ablation study, including the use of δ -coordinates, RBF-weighted function as well as the implicit surface encoder, as highlighted in Figure 2. We use PointNet, where only per-point MLP layers are applied on each point in the point cloud, and DGCNN, which dynamically build graphs according to K nearest neighbors, as the baseline. RBF-weighted δ -coordinates (as described in Equation (5)), uniformly weighted δ -coordinates (described

in Equation (4)), and implicit surface encoder shown at the top of Figure 2 are considered as subjects of ablation studies. We observe that our fine-grained geometric feature representation in Equation (6) results in larger reduction in dB error, especially at higher frequencies. This demonstrates improved accuracy of our approach over [Tang *et al.*, 2020]. Moreover, our discrete Laplacian based representation is more robust to the Gaussian noise in the test set compared to DGCNN. We visualize some results from these networks in Figure 3 and in Appendix.

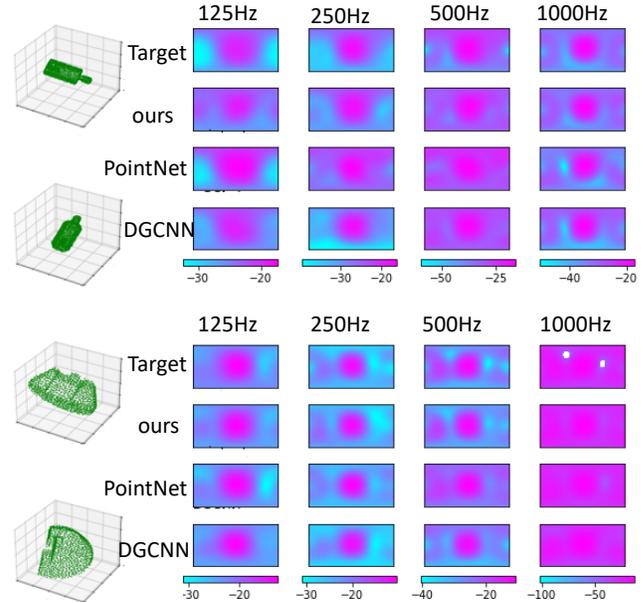


Figure 3: **ASF predictions and comparison:** We expand the predicted spherical harmonics coefficients onto a latitude-longitude (lat-long) map, representing the directional distribution of the scattering field. Here we show 2 sample point clouds from the unseen test set. Each point cloud is visualized from two viewing angles, followed by their groundtruth/target (Row 1), prediction using our method (Row 2), PointNet (Row 3), and DGCNN (Row 4). More results visualizations are in Appendix.

5.3 Application to Interactive Sound Propagation

We use our ASF computation in a hybrid sound propagation system. Since our learning method can compute the ASF in less than 1ms, our approach involves no precomputation and we are able to handle deforming objects or changing topologies. In particular, we combine the ASFs with state-of-the-art ray-tracing based acoustic simulator [Schissler and Manocha, 2017]. In our hybrid system, the learned ASFs are used to simulate wave acoustics up to $1000Hz$ to compensate the low-frequency effects like sound diffraction and interference. This choice of the simulation frequency of $1000Hz$ is similar to [Raghuvanshi and Snyder, 2018] and [Rungta *et al.*, 2018], as it can provide plausible sounds for games, virtual reality, and computer-aided designs.

Conventional ray-tracing algorithm finds direct and indirect ray paths from a sound source to a receiver/listener. Next, a time delay and energy damping is calculated

²<https://shools.oca.eu/shools/public/index.html>

Methods {Eq. 5 — surface encoder}	Error in dB				#Param
	125Hz	250Hz	500Hz	1000Hz	
{ \mathcal{X} — \mathcal{X} }	3.49	3.56	3.71	4.23	30k
{ \checkmark — \mathcal{X} }	3.38	3.41	3.57	4.47	30k
{ \mathcal{X} — \checkmark }	3.28	3.38	3.52	3.85	30k
{ \checkmark — \checkmark } (ours)	2.98	3.06	3.22	3.76	34k
PointNet [Charles <i>et al.</i> , 2017]	3.66	3.44	3.43	4.16	80k
DGCNN [Wang <i>et al.</i> , 2019]	3.71	3.56	3.59	4.21	103k

Table 1: **Ablation study:** In this evaluation, we compare the performance on uniform δ - coordinate in Equation (4), weighted δ - coordinates in Equation (5) and the surface encoder in Equation (6) on our test dataset (including 10k objects) at frequency bands {125Hz, 250Hz, 500Hz, 1000Hz}. The best result for each frequency is highlighted in **bold** (lower error is better). We alter between choosing Equation (4) and Equation (5) that results in four different combinations: Row 1 and 2, 3 and 4. Next, we experiment the use of implicit surface encoder (Row 3 and 4). Our proposed network design (Row 4) shows superior performance in terms of ASF approximation for all four frequencies.

Methods	Error in dB			
	125Hz	250Hz	500Hz	1000Hz
PointNet	5.88	5.50	6.02	8.37
DGCNN	7.91	6.65	7.31	11.32
Ours	5.20	5.09	4.70	5.51

Table 2: **Robustness test:** In this robustness test, we add i.i.d. noise $\sim N(0, 0.05)$ to every point in the point cloud in the test set. Note that DGCNN finds K - neighbors at each iteration, similar to our approach. However, our algorithm is more robust to the noisy data input thanks to the design of the shape Laplacian based point representation. PointNet was used in [Tang *et al.*, 2020] and our approach is more accurate at higher frequencies.

according to the path length and the order of reflections of each path. This information is used to compose an impulse response for real-time sound rendering [Kuttruff, 2016].

The core of the ray-tracing process is to sample a large number of reflection directions and their corresponding energy when a ray hits a surface. A pure geometric sound propagation system assumes specular and diffuse reflections based on high-frequency approximation (Figure 5 left). We use the ASF to approximate the diffraction and other wave-effects. Specifically, when an incoming ray hits a sound scatterer, we sample the outgoing directions among all directions and calculate the energy decay by evaluating the ASF at each direction (θ_i, ϕ_i) for frequency ω (Figure 5 right). This hybrid wave-ray formulation is general enough to be used with most sound propagation engines.

We also evaluate the sound rendering in the four benchmark scenes with different complexity, as shown in Figure 4. We demonstrate that our ASFs can generate smooth and realistic sound renderings (in supplemental video).

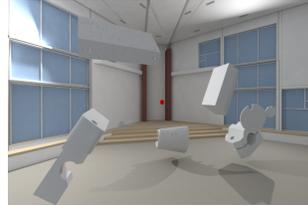
6 Conclusion and Limitations

We present a novel geometric deep learning approach that infers the 3D acoustic scattering field induced by an object from its shape. Our algorithm uses discrete-laplacian based implicit function encoders to accurately capture the geometric properties of an object and achieves lower error compared

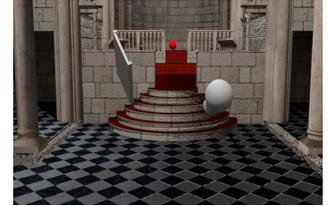


(a) Floor: One static sound scatterer in open space. We compute the ASF and perform raytracing at 10.65ms/frame.

(b) Havana: Two moving walls in half-open space. We compute the ASF for each position of the wall for sound propagation at 6.78ms/frame.



(c) Trinity: Six flying objects in a large indoor room. We compute the ASF for each object and use that for sound propagation at 12.95ms/frame.



(d) Sibenik: Two disjoint revolving objects in a church. We compute their ASF and use them for sound propagation at 6.87ms/frame.

Figure 4: Benchmark scenes used for audio-visual rendering in our supplemental video. These are dynamic scenes where objects come in close proximity and change topologies. Our learning methods can compute accurate scattering fields and combine with ray tracing for interactive sound propagation.

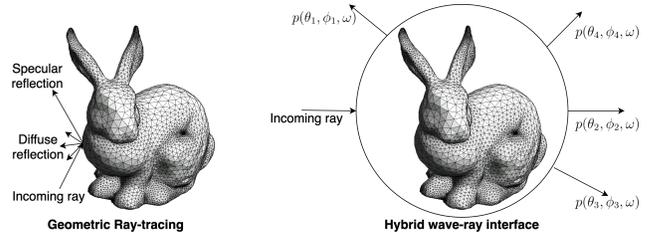


Figure 5: Pure geometric ray-tracing (left) vs our hybrid wave-ray coupling (right). Our learned ASF is used to approximate the scattering effects and compute multiple output directions.

with existing general point-based learning methods. We use these scattered fields for interactive sound rendering and highlight their application in complex benchmarks.

Our current implementation is limited to sound-hard surfaces and a limited range of simulation frequencies, though these are not inherent limitations of our method. The simulation time needed to create a large training set can grow as a cubic function of the sound frequency. Incorporating more types of sound surfaces can further increase the training overhead of all geometric deep learning methods.

Acknowledgments

This work was supported in part by ARO grants W911NF-18-1-0313 and W911NF-19-1-0315, NSF grants #1910940, Adobe, Facebook and Intel.

References

- [Bianco *et al.*, 2019] Michael J Bianco, Peter Gerstoft, James Traer, Emma Ozanich, Marie A Roch, Sharon Gannot, and Charles-Alban Deledalle. Machine learning in acoustics: Theory and applications. *The Journal of the Acoustical Society of America*, 146(5):3590–3628, 2019.
- [Charles *et al.*, 2017] R. Qi Charles, Hao Su, Mo Kaichun, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *IEEE CVPR*, Jul 2017.
- [Fan *et al.*, 2020] Ziqi Fan, Vibhav Vineet, Hannes Gamper, and Nikunj Raghuvanshi. Fast acoustic scattering using convolutional neural networks. In *ICASSP*, pages 171–175. IEEE, 2020.
- [Hanocka *et al.*, 2019] Rana Hanocka, Amir Hertz, Noa Fish, Raja Giryes, Shachar Fleishman, and Daniel Cohen-Or. Meshcnn: a network with an edge. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019.
- [Heath and Sudderth, 1976] David Heath and William Sudderth. De finetti’s theorem on exchangeable variables. *The American Statistician*, 30(4):188–189, 1976.
- [Koch *et al.*, 2019] Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. Abc: A big cad model dataset for geometric deep learning. In *IEEE CVPR*, June 2019.
- [Kuttruff, 2016] Heinrich Kuttruff. *Room Acoustics*. Taylor & Francis Group, London, U. K., 6th edition, 2016.
- [Kůrková, 1992] Věra Kůrková. Kolmogorov’s theorem and multilayer neural networks. *Neural Networks*, 5(3):501–506, 1992.
- [Liu and Manocha, 2020] Shiguang Liu and Dinesh Manocha. Sound synthesis, propagation, and rendering: A survey. *arXiv preprint arXiv:2011.05538*, 2020.
- [Liu *et al.*, 2019] Shichen Liu, Shunsuke Saito, Weikai Chen, and Hao Li. Learning to infer implicit surfaces without 3d supervision. In *Advances in Neural Information Processing Systems*, pages 8293–8304, 2019.
- [Mo *et al.*, 2019] Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy J. Mitra, and Leonidas J. Guibas. Structurenet. *ACM Transactions on Graphics (TOG)*, 38(6):1–19, Nov 2019.
- [Park *et al.*, 2019] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *IEEE CVPR*, pages 165–174, 2019.
- [Pulkki and Svensson, 2019] Ville Pulkki and U Peter Svensson. Machine-learning-based estimation and rendering of scattering in virtual reality. *The Journal of the Acoustical Society of America*, 145(4):2664–2676, 2019.
- [Qi *et al.*, 2017] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108, 2017.
- [Raghuvanshi and Snyder, 2018] Nikunj Raghuvanshi and John Snyder. Parametric directional coding for precomputed sound propagation. *ACM Transactions on Graphics (TOG)*, 37(4):108, 2018.
- [Rumelhart *et al.*, 1985] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- [Rungta *et al.*, 2018] Atul Rungta, Carl Schissler, Nicholas Rewkowski, Ravish Mehra, and Dinesh Manocha. Diffraction kernels for interactive sound propagation in dynamic environments. *IEEE Transactions on Visualization and Computer Graphics*, 24(4):1613–1622, 2018.
- [Schissler and Manocha, 2017] Carl Schissler and Dinesh Manocha. Interactive sound propagation and rendering for large multi-source scenes. *ACM Transactions on Graphics (TOG)*, 36(1):2, 2017.
- [Smirnov *et al.*, 2020] Dmitriy Smirnov, Matthew Fisher, Vladimir G Kim, Richard Zhang, and Justin Solomon. Deep parametric shape predictions using distance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 561–570, 2020.
- [Sorkine, 2006] Olga Sorkine. Differential representations for mesh processing. In *Computer Graphics Forum*, volume 25, pages 789–807. Wiley Online Library, 2006.
- [Tang *et al.*, 2020] Zhenyu Tang, Hsien-Yu Meng, and Dinesh Manocha. Learning acoustic scattering fields for dynamic interactive sound propagation. *arXiv preprint arXiv:2010.04865*, 2020.
- [Wang *et al.*, 2019] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics*, 38(5):1–12, Oct 2019.
- [Xu *et al.*, 2019] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. In *Advances in Neural Information Processing Systems*, pages 490–500, 2019.
- [Zhou and Tuzel, 2018] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. *IEEE CVPR*, Jun 2018.