

Image2Reverb: Cross-Modal Reverb Impulse Response Synthesis

Nikhil Singh
MIT
Media Lab

nsingh1@mit.edu

Jeff Mentch
Harvard University
SHBT

jsmentch@g.harvard.edu

Jerry Ng
MIT
Mechanical Engineering

jerryng@mit.edu

Matthew Beveridge
MIT
EECS

mattbev@mit.edu

Iddo Drori
MIT
EECS

idrori@mit.edu



Figure 1. An application of our work: matching speaker audio to expected reverberation of a video conference virtual background. Project page with examples is available at <https://web.media.mit.edu/~nsingh1/image2reverb>.

Abstract

Measuring the acoustic characteristics of a space is often done by capturing its impulse response (IR), a representation of how a full-range stimulus sound excites it. This work generates an IR from a single image, which can then be applied to other signals using convolution, simulating the reverberant characteristics of the space shown in the image. Recording these IRs is both time-intensive and expensive, and often infeasible for inaccessible locations. We use an end-to-end neural network architecture to generate plausible audio impulse responses from single images of acoustic environments. We evaluate our method both by comparisons to ground truth data and by human expert evaluation. We demonstrate our approach by generating plausible impulse responses from diverse settings and formats including well known places, musical halls, rooms in paintings, images from animations and computer games, synthetic environments generated from text, panoramic images, and video conference backgrounds.

1. Introduction

An effective and widely used method of simulating acoustic spaces relies on audio impulse responses (IRs) and

convolution [38, 31]. Audio IRs are recorded measurements of how an environment responds to an acoustic stimulus. IRs can be measured by recording a space during a burst of white noise like a clap, a balloon pop, or a sinusoid swept across the range of human hearing [28]. Accurately capturing these room impulse responses requires time, specialized equipment, knowledge, and planning. Directly recording these measurements may be entirely infeasible in continuously inhabited or inaccessible spaces of interest. End-to-end IR estimation has far ranging applications relevant to fields including music production, speech processing, and generating immersive extended reality environments. Our Image2Reverb system directly synthesizes IRs from images of acoustic environments. This approach removes the barriers to entry, namely cost and time, opening the door for a broad range of applications.

In this work we model IR generation as a cross-modal paired-example domain adaptation problem and apply a conditional GAN [10, 12, 23] to synthesize plausible audio impulse responses conditioned on images of spaces. Next, we will describe related work that informs our approach.

2. Related Work

Artificial reverberation. Historically, recording studios built reverberant chambers with speakers and microphones

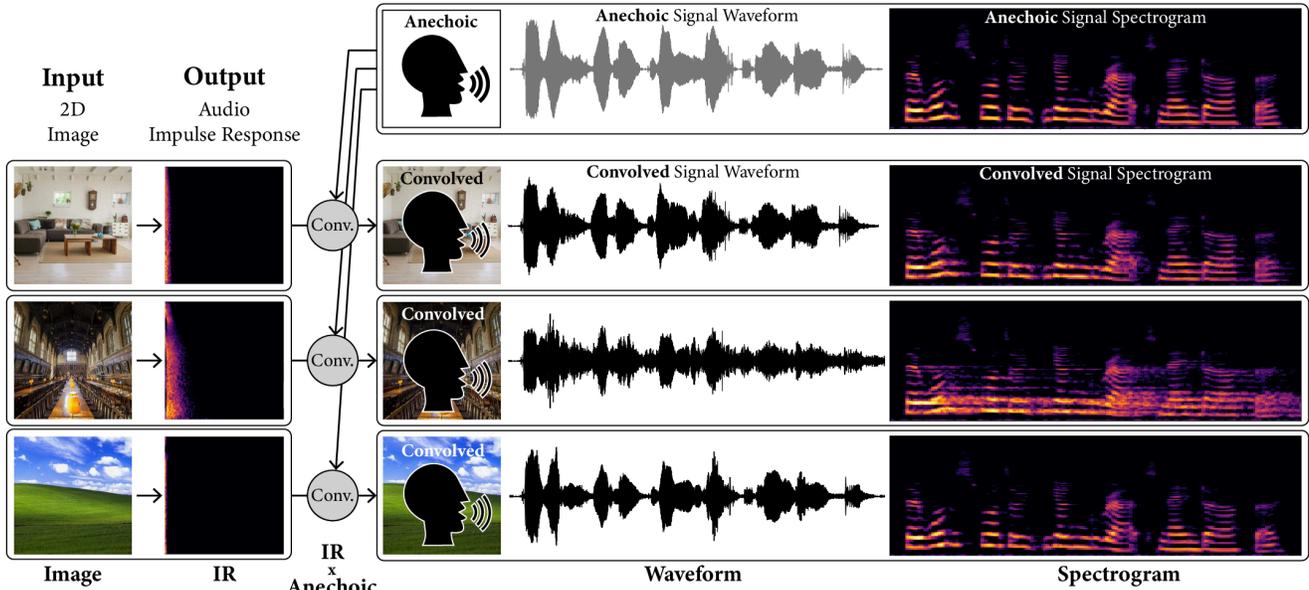


Figure 2. Generating audio impulse responses from images. Left: given an image of an acoustic environment as input, our model generates the corresponding audio impulse response as output. Right: generated impulse responses are convolved with an anechoic (free from echo) audio recording making that recording sound as if it were in the corresponding space. Waveforms and spectrograms are shown of the source anechoic signal and the same signal after convolution with the corresponding synthesized IR. All spectrograms are presented on a mel scale. Image2Reverb is the first system demonstrating end-to-end synthesis of realistic IRs from single images.

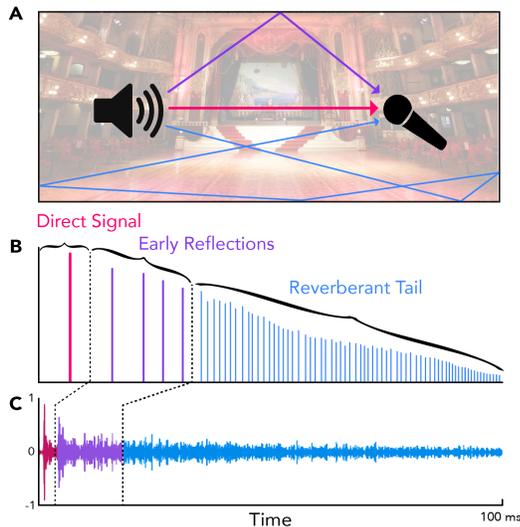


Figure 3. Impulse response overview. (A) Sound waves propagate across multiple paths as they interact with and reflect off their environment. These paths include the direct path from source to listener, early reflections including 1st and higher order reflections (after reflecting off 1 or more surfaces) and a more diffuse tail as they trail off and become more densely packed in time. These reflections make up the impulse response of the environment illustrated (B) schematically and (C) as a waveform.

to apply reverb to pre-recorded audio directly within a physical space [30]. Reverberation circuits, first proposed in the

1960s, use a network of filters and delay lines to mimic a reverberant space [34]. Later, Digital algorithmic approaches applied numerical methods to simulate similar effects. Conversely, convolution reverb relies on audio recordings of a space’s response to a broadband stimulus, typically a noise burst or sine sweep. This results in a digital replica of a space’s reverberant characteristics, which can then be applied to any audio signal [1].

Convolutional neural networks have been used for estimating late-reverberation statistics from images [16, 17], though not to model the full audio impulse response from an image. This work is based on the finding that experienced acoustic engineers readily estimate a space’s IR or reverberant characteristics from an image [15]. Room geometry has also been estimated from 360-degree images of four specific rooms [29], and used to create virtual acoustic environments which are compared with ground-truth recordings, though again IRs are not directly synthesized from the images. A related line of work synthesizes spatial audio based on visual information [19, 8, 14]. Prior work exists on synthesis of IRs using RNNs [32], autoencoders [36], and GANs: IR-GAN [27] uses parameters from real world IRs to generate new synthetic IRs; whereas our work synthesizes an audio impulse response directly from an image.

Generative models for audio. Recent work has shown that GANs are amenable to audio generation and can result in more globally coherent outputs [6]. GANSynth [7]

generates an audio sequence in parallel via a progressive GAN architecture allowing faster than real-time synthesis and higher efficiency than the autoregressive WaveNet [39] architecture. Unlike WaveNet which uses a time-distributed latent coding, GANSynth synthesizes an entire audio segment from a single latent vector. Given our need for global structure, we create a fixed-length representation of our input and adapt our generator model from this approach.

Measured IRs have been approximated with shaped noise [18, 3]. While room IRs exhibit statistical regularities [37] that can be modeled stochastically, the domain of this modeling is time and frequency limited [2], and may not reflect all characteristics of real-world recorded IRs. Simulating reverb with ray tracing is possible but prohibitively expensive for typical applications [33]. By directly approximating measured audio IRs at the spectrogram level, our outputs are immediately applicable to tasks such as convolution reverb, which applies the reverberant characteristics of the IR to another audio signal.

Cross-modal translation. Between visual and auditory domains, conditional GANs have been used for translating between images and audio samples of people playing instruments [4]. Our work builds on this by applying state-of-the-art architectural approaches for scene analysis and high quality audio synthesis, tuned for our purposes.

3. Methods

Here we describe the dataset, model, and algorithm.

3.1. Dataset

Data aggregation. We curated a dataset of 265 different spaces totalling 1169 images and 738 IRs. From these, we produced a total of 11234 paired examples with a train-validation-test split of 9743-154-1957. These are assembled from sources including the OpenAIR dataset [24], other libraries available online, and web scraping. Many examples amount to weak supervision, due to the low availability of data: for example, we may have a “kitchen” impulse response without an image of the kitchen in which it was recorded. In this case, we augmented with plausible kitchen scenes, judged by the researchers, gathered via web scraping and manual filtering. Although this dataset contains high variability in several reverberant parameters, e.g. early reflections and source-microphone distance, it allows us to learn characteristics of late-field reverberation.

Data preprocessing. Images needed to be filtered manually to remove duplicates, mismatches such as external pictures of an indoor space, examples with significant occlusive “clutter” or excessive foreground activity, and intrusive watermarks. We then normalized, center-cropped at the

max width or height possible, and downsampled to 224x224 pixels. We converted the audio IR files to monaural signals; in the case of Ambisonic B-Format sources we extracted the W (omnidirectional) channel, and for stereo sources we computed the arithmetic mean of channels. In some cases, 360-degree images were available and in these instances we extract rectilinear projections, bringing them in line with the standard 2D images in our dataset.

Audio representation. Our audio representation is a log magnitude spectrogram. We first resampled the audio files to 22.050kHz and truncate them to 5.94s in duration. This is sufficient to capture general structure and estimate reverberant characteristics for most examples. We then apply a short-time Fourier transform with window size ($M = 1024$) and hop size ($R = 256$), before trimming the Nyquist bin, resulting in square 512x512 spectrograms. Finally, we take $\log(|X|)$ where $|X|$ represents the magnitude spectrogram; audio IRs typically contain uncorrelated phase, which does not offer structure we can replicate based on the magnitude.

3.2. Model

Components. Our model employs a conditional GAN with an image encoder that takes images as input and produces spectrograms. This overall design, with an encoder, generator, and conditional discriminator, is similar to that which Mentzer et al. [22] applied to obtain state-of-the-art results on image compression, among many other applications. The generator and discriminator are deep convolutional networks based on the GANSynth [7] model (non-progressive variant), with modifications to suit our dataset, dimensions, and training procedure.

The encoder module combines image feature extraction with depth estimation to produce latent vectors from two-dimensional images of scenes. For depth estimation, we use the pretrained Monodepth2 network [9], a monocular depth-estimation encoder-decoder network which produces a one-channel depth map corresponding to our input image. The main feature extractor is a ResNet50 [13] pretrained on Places365 [41] which takes a four-channel representation of our scene including the depth channel (4x224x224). We add randomly initialized weights to accommodate the additional input channel for the depth map. Since we are fine-tuning the entire network, albeit at a low learning rate, we expect it will learn the relevant features during optimization. Our architecture’s components are shown in Figure 4.

Objectives. We use the least-squares GAN formulation (LSGAN) [20]. For the discriminator:

$$\begin{aligned} \min_D V(D) = & \mathbb{E}_{\mathbf{y} \sim p_{data}(\mathbf{y})} [(1 - D(\mathbf{y} | E(\tilde{\mathbf{x}})))^2] \\ & + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [(D(G(\mathbf{z}) | E(\tilde{\mathbf{x}})))^2] \end{aligned} \quad (1)$$

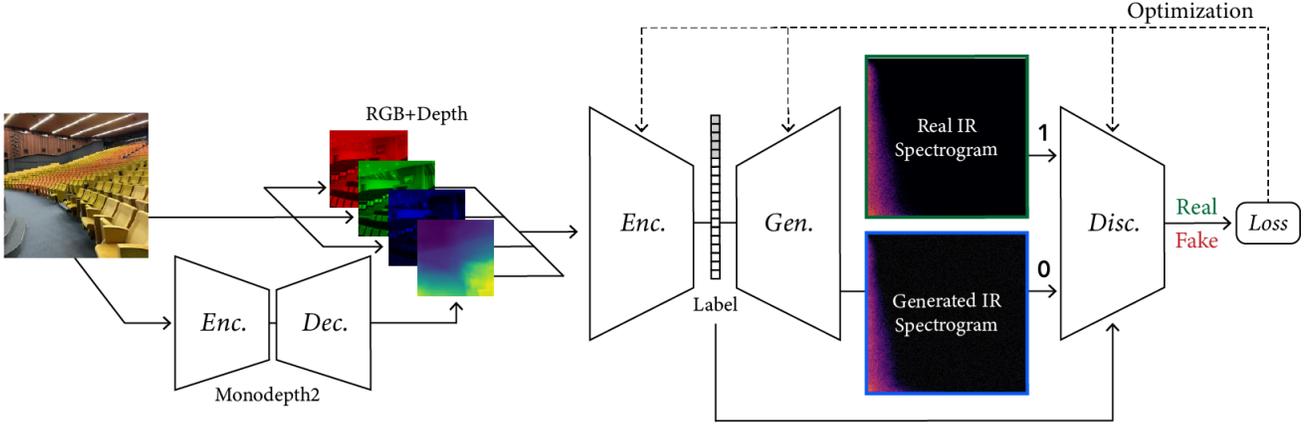


Figure 4. System architecture. Our system consists of autoencoder and GAN networks. Left: An input image is converted into 4 channels: red, green, blue and depth. The depth map is estimated by Monodepth2, a pre-trained encoder-decoder network. Right: Our model employs a conditional GAN. An image feature encoder is given the RGB and depth images and produces part of the Generator’s latent vector which is then concatenated with noise. The Discriminator applies the image latent vector label at an intermediate stage via concatenation to make a conditional real/fake prediction, calculating loss and optimizing the Encoder, Generator, and Discriminator.

Algorithm 1: Forward and backward passes through the Image2Reverb model. Notation is explained in Table 1.

Input:

Monodepth2: $x \sim X$; Encoder $\tilde{x} \sim \tilde{X}$; Generator: $z = E(\tilde{x}) \oplus u$; Discriminator: $(G(z), E(\tilde{x}))$ OR $(y, E(\tilde{x}))$;
Parameters: (weight variables);

Output:

Monodepth2: x_d ; Encoder: $E(\tilde{x})$; Generator: $G(z)$;
Discriminator: $D(G(z), E(\tilde{x}))$ OR $D(y, E(\tilde{x}))$;

for number of epochs do

Sample B training images;
Get depth $x_d = M(x)$;
Append depth features to RGB channels $(y \oplus y_d)$;
Encoder image to feature-vector $(E(\tilde{x}))$;
Append noise $(z = E(\tilde{x}) \oplus u)$;
Generate spectrogram $(G(z))$;
Forward pass through discriminator with either fake or real spectrogram $(D(G(z)|E(\tilde{x}))$ OR $D(y|E(\tilde{x})))$;
Backward pass: update parameters for discriminator (W_D), generator (W_G), and encoder (W_E);

end

Notation	Definition
x	input image
x_d	estimated depth map
\oplus	concatenation operator
\tilde{x}	image with depth map $(x \oplus x_d)$
y	Real spectrogram
E, G, D	Encoder, Generator, Discriminator
M	Monodepth2 Encoder-Decoder
W_*	weights for a model
u	Noise, $u \sim \mathcal{N}(0, 1)$
z	Latent vector, encoder output and noise $(E(\tilde{x}) \oplus u)$

Table 1. Notation and definitions for variables indicated in different parts of this paper.

image and audio settings. Second, we introduce a domain-specific term that performs an estimation of the T_{60} values, the time it takes for the reverberation to decay by 60dB, for the real and generated samples, and returns the absolute percent error between the two scaled by a hyperparameter ($\lambda_b = 100$ again). We term the differentiable T_{60} proxy measure T_{60_p} . To compute this for log-spectrogram x , we first get the linear spectrogram e^x and then sum along the time axis to obtain a fullband amplitude envelope. We use Schroeder’s backward integration method to obtain a decay curve from the squared signal, and linearly extrapolate from the $-20dB$ point to get a T_{60} estimate. In all:

$$\begin{aligned} \min_G V(G) &= \mathbb{E}_{z \sim p_z(z)} [(1 - D(G(z) | E(\tilde{x})))^2 \\ &+ \lambda_a \|G(z) - y\|_1 \\ &+ \lambda_b \left| \frac{T_{60_p}(G(z)) - T_{60_p}(y)}{T_{60_p}(y)} \right|] \end{aligned} \quad (2)$$

For the generator, we introduce two additional terms to encourage realistic and high-quality output. First, we add an ℓ_1 reconstruction term, scaled by a hyperparameter ($\lambda_a = 100$ in our case). This is a common approach in

Training. We train our model on 8 NVIDIA 1080 Ti GPUs. Three Adam optimizers for each of the Generator, Discriminator, and Encoder were used to optimize the networks’ parameter weights. Hyperparameters are noted in Table 2. We make our model and code publicly available ¹.

Parameter	Value
η_G	4e-4
η_D	2e-4
η_E	1e-5
β	(0.0, 0.99)
ϵ	1e-8

Table 2. Hyperparameters for the Generator, Discriminator, and Encoder initial learning rates, the optimizer beta (β), and epsilon (ϵ) for the Adam optimizers we use (one each for D, G, E)

4. Results

Using Image2Reverb we are able to generate perceptually plausible impulse responses for a diverse set of environments. In this section, we provide input-output examples to demonstrate the capabilities and applications of our model and also review results of a multi-stage evaluation integrating domain-specific quantitative metrics and expert ratings. Our goal is to examine output quality and conditional consistency, generally considered important for conditional GANs [5] and most relevant for our application.

4.1. Examples

We present several collections consisting of diverse examples in our supplementary material, with inputs curated to illustrate a range of settings of interest including famous spaces, musical environments, and entirely virtual spaces. All examples are made available as audiovisual collections² and were generated with a model trained in around 12 hours, with 200 epochs on a virtual machine. Figure 5 shows examples from our test set that were used in our expert evaluation (4 of 8, one from each category of: Small, Medium, Large, and Outdoor). We convolve a spoken word anechoic signal with the generated IRs for the reader to hear. Figure 6 takes images of diverse scenes (art, animation, historical/recognizable places) as inputs. Figure 7 demonstrates how sections of 360-degree equirectangular images are cropped, projected, and passed through our model to generate IRs of spaces for immersive VR environments.

We strongly encourage the reader to explore these examples on the accompanying web page. We include examples of musical performance spaces, artistic depictions (drawings, paintings), 3D animation scenes, synthetic images from OpenAI’s DALL•E, as well as real-world settings that present challenges (e.g. illusions painted on walls, reflections, etc.). These are largely created with real-world

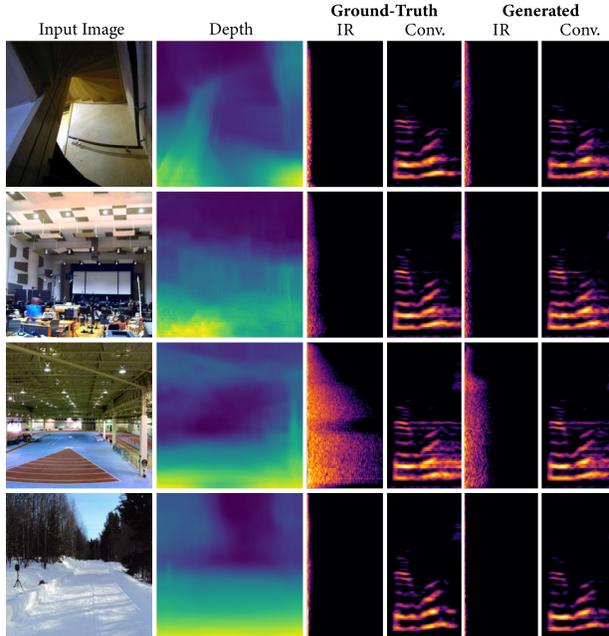


Figure 5. Ground-truth measured IRs vs generated IRs. Columns show input images, depth maps, measured IRs with corresponding convolved speech, and generated IRs with corresponding convolved speech. Larger indoor spaces here tend to exhibit greater T_{60} times with longer measured impulse responses. The outdoor scene has a very short measured IR and corresponding generated IR. Input images are all examples that were used in the expert survey and were drawn from the test set.

environments for which we may not have ground truth IRs, demonstrating how familiar and unusual scenes can be transformed in this way.

4.2. Ablation Study

To understand the contribution of key architectural components and decisions, we perform a study to characterize how removing each affects test set T_{60} estimation after 50 training epochs. The three components are the depth maps, the $T_{60,p}$ objective term, and the pretrained Places365 weights for the ResNet50 encoder. Figure 8 shows T_{60} error distributions over the test set for each of these model variants, and Table 3 reports descriptive statistics.

Our model reflects better mean error (closer to 0%) and less dispersion (a lower standard deviation) than the other variants. The former is well within the just noticeable difference (JND) bounds for T_{60} , often estimated as being around 25-30% for a musical signal [21]. Additionally, this is an upper bound on authenticity: a more rigorous goal than perceptual plausibility [26]. The lower standard deviation indicates generally more consistent performance from this model across different examples, even in the presence of some that cause relatively large estimation errors due to in-

¹Model and code: <https://github.com/nikhil Singhmus/image2reverb>

²Audiovisual samples: <https://web.media.mit.edu/~nsingh1/image2reverb/>

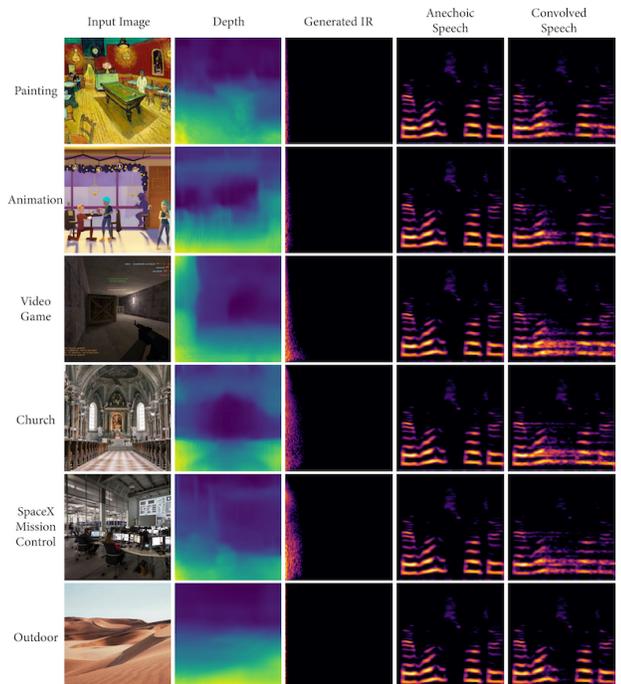


Figure 6. Generated IR examples. Columns show input images, depth maps, generated IRs, and a dry anechoic speech signal before and after the generated IR was applied via convolution. Input images come from a variety of spaces which illustrate possible applications of our model. Some images are synthetic, including: an oil painting, a 3D animation still, and a video game screenshot. Others come from real-world scenes like a church (where music is often heard), a famous yet inaccessible space (SpaceX), and an outdoor desert scene. Larger indoor spaces tend to exhibit longer impulse responses as seen here.



Figure 7. VR. Impulse responses generated from an equirectangular 360-degree image by sampling points on a sphere, cropping and applying a rectilinear projection to the resulting image, and feeding them into our model. This demonstrates how our model directly generates realistic impulse responses of panoramic virtual reality compatible images. Future work may allow generation of impulse responses using an entire 360-degree image, though at present there is a lack of paired data available for training.

correct interpretation of relevant qualities in the image, or inaccurate/noisy synthesis or estimation.

4.3. Expert Evaluation

Following the finding that experienced acoustic engineers readily estimate a space’s reverberant characteristics from an image [15], we designed an experiment to evalu-

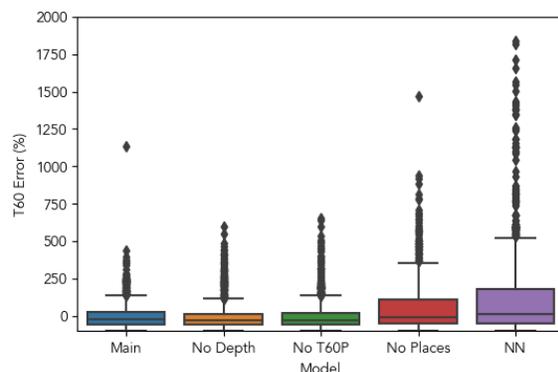


Figure 8. T_{60} estimation error (%) distributions from each model version. T_{60} estimates how long it takes the reverberation to decay by $60dB$. “Main” is our architecture as described earlier, “No Depth” omits the depth maps, “No T60P” omits the differentiable T_{60p} objective term, and “No Places” uses randomly initialized encoder weights. “NN” applies a nearest-neighbor approach with Places365-ResNet50 embeddings for images (errors clipped to 2000% for clarity). Descriptive statistics are given in Table 3.

		Main	-Depth	$-T_{60p}$	-P365	NN
T_{60} Err (%)	μ	-6.03	-9.17	-7.1	43.15	149
	σ	78.8	83.1	85.97	144.3	491.02

Table 3. T_{60} estimation error (%) statistics from each model version. “Main” is our architecture as described earlier, “-Depth” omits depth maps, “ $-T_{60p}$ ” omits the differentiable T_{60p} objective term, and “-P365” does not use the pretrained Places365 weights for the ResNet50 encoder. “NN” indicates a nearest-neighbor approach with Places365-ResNet50 embeddings for images. For mean and median, values closer to 0 reflect better performance. For the standard deviation, lower values reflect better performance. Distributions are visualized in Figure 8.

ate our results. We note that this experiment is designed to estimate comparative perceptual plausibility, rather than (physical) authenticity (e.g. by side-by-side comparison to assess whether any difference can be heard). These goals have been differentiated in prior work [26]. We selected two arbitrary examples from each of the four scene categories and recruited a panel of 31 experts, defined as those with significant audio experience, to participate in a within-subjects study. For each of these examples, we convolved an arbitrary anechoic signal with the output IR, as well as the ground truth IR. These 16 samples were presented in randomized order and participants were instructed to rate each on a scale from 1 to 5 based on 1) reverberation quality, and 2) realism or “match” between their expected reverb based on the image and the presented signal with reverb applied. Participants answered one reverb-related screening question to demonstrate eligibility, and two attention check questions at the end of the survey. The four scene categories

Rating	Scene	DoF	p
Quality	Large	56	< .001
Quality	Medium	56	.28
Quality	Outdoor	56	.62
Quality	Small	56	< .001
Match	Large	56	< .001
Match	Medium	56	.006
Match	Outdoor	56	.29
Match	Small	56	< .05

Table 4. Simple main effect tests for equivalence between real and generated IRs across different categories of scenes. We use paired two one-sided tests with bounds (ϵ) of 1 and Bonferroni-adjusted p -values. These results suggest that real vs. fake ratings are statistically equivalent within one rating unit (the resolution of the rating scale) for large and small quality ratings, and large, medium, and small match ratings. Notably, outdoor scenes contribute to the difference between real and fake IRs and medium-sized scenes contribute to differences in quality.

are: Large, Medium, Outdoor, and Small. These demonstrate diversity in visual-reverb relationships. The dependent variables are quality and match ratings, and the independent variables are IR source (real vs. fake) and scene category (the four options listed previously). We first test our data for normality with D’Agostino and Pearson’s omnibus test [25], which indicates that our data is statistically normal ($p > .05$).

A two-way repeated-measures ANOVA revealed a statistically significant interaction between IR source and scene category for both quality ratings, $F(3, 90) = 7.04$, $p \leq .001$, and match ratings, $F(3, 90) = 3.73$, $p = .02$ (reported p -values are adjusted with the Greenhouse-Geisser correction [11]). This indicates that statistically significant differences between ratings for real and fake IR reverbs depend on the scene category. Per-participant ratings and rating changes, overall and by scene, are shown in Figure 9.

Subsequent tests for simple main effects with paired two one-sided tests indicate that real vs. fake ratings are statistically equivalent ($p < .05$) for large and small quality ratings, and large, medium, and small match ratings. These tests are carried out with an ϵ of 1 (testing for whether the means of the two populations differ by at least 1). Results are shown in Table 4. Notably, outdoor scenes appear to contribute to the rating differences between real and fake IRs. We conjecture this is due to outdoor scenes being too different a regime from the vast majority of our data, which are indoor, to model effectively. Additionally, medium-sized scenes appear to contribute to differences in quality.

4.4. Model Behavior and Interpretation

Effect of varying depth. We compare the full estimated depth map with constant depth maps filled with either 0 or 0.5 (chosen based on the approximate lower and upper bounds of our data). We survey the distributions of gener-

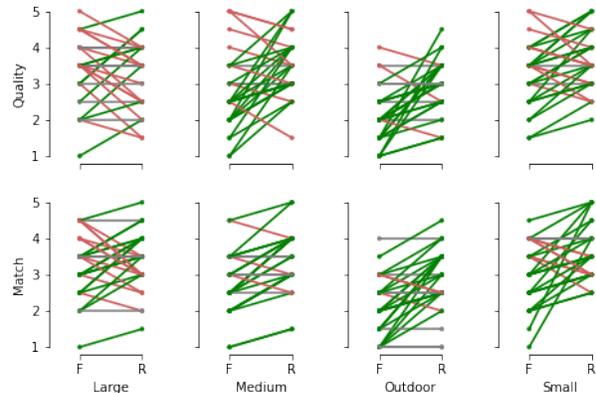


Figure 9. Expert evaluation results. Paired plots showing per-participant quality and match differences in rating for each scene category. Green lines indicate higher rating for real IRs, red lines for generated IRs, and grey lines equivalent ratings for both.

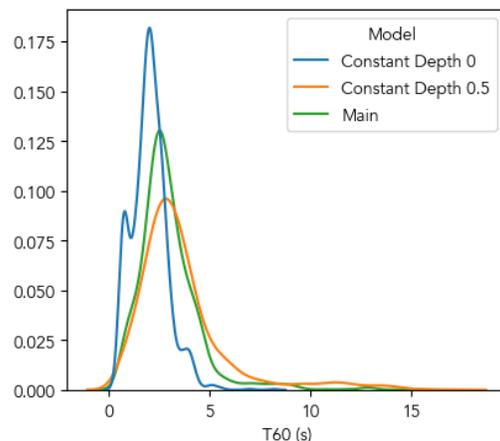


Figure 10. Effect of Depth on T_{60} . Distributions of estimated T_{60} values for the model with estimated depth maps, plus constant depth maps set to either 0 (low) or 0.5 (high). Manipulating the depth value allows us to “suggest” smaller or larger scenes, i.e. bias the output of the model. Table 5 shows corresponding descriptive statistics. These results indicate a level of “steerability” for the model’s behavior in human-in-the-loop settings.

ated IRs’ T_{60} values over our test set, the results of which are shown in Figure 10. Table 5 reports descriptive statistics for these distributions, showing that the main model’s output IRs’ decay times are biased lower by the 0-depth input and higher by the 0.5-depth input respectively. These may indicate some potential for steering the model in interactive settings. We do note, however, that behavior with constant depth values greater than 0.5 is less predictable. This may be due to the presence of outdoor scenes, for which the scene’s depth may not be correlated with IR duration.

Effect of transfer learning. To understand which visual features are important to our encoder, we use Gradient-

	Main	Depth 0	Depth 0.5
μ	2.07	2.01	3.62
σ	1.54	0.87	2.36
$Mdn.$	2.69	2.00	3.07

Table 5. Descriptive statistics for the model with estimated depth maps, as well as constant depth maps set to either 0 or 0.5. The full depth map’s results are between that of the 0 and 0.5 depth maps. Figure 10 visualizes the corresponding distributions.

weighted Class Activation Mapping (Grad-CAM) [35]. Grad-CAM is a popularly applied strategy for visually interpreting convolutional neural networks by localizing important regions contributing to a given target feature (or class in a classification setting). We produce such maps for our test images with both the ResNet50 pre-trained on Places365 dataset, as well as the final encoder model. All resulting pairs exhibit noticeable differences; we check for this with the structural similarity index (SSIM) metric [40], which is below 0.98 for all examples.

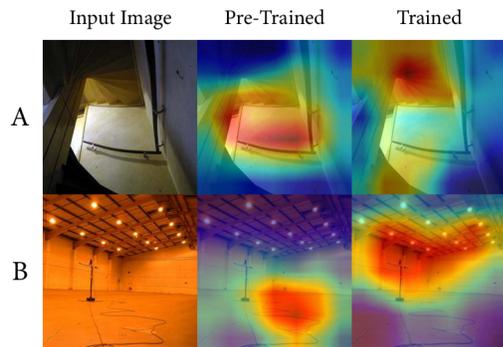


Figure 11. Grad-CAMs for images passed through the pre-trained Places365 ResNet50 encoder vs. our fine-tuned encoder, showing movement towards significant reflective areas for (A) a small, and (B) a large environment. The fine-tuned model’s activations highlight larger reflective surfaces: depth of staircase for (A) vs. railing that may be more optimal for scene identification, and wall-to-ceiling corner plus surrounding areas for (B).

We qualitatively survey these and identify two broad change regimes, which are illustrated with particular examples. First, we observe that the greatest-valued feature is often associated with activations of visual regions corresponding to large reflective surfaces. Examples are shown in Figure 11. Often, these are walls ceilings, windows, and other surfaces in reflective environments. Second, we find that textured areas are highlighted in less reflective environments. Examples of these are shown in Figure 12. These may correspond to sparser reflections and diffusion.

Limitations and future work. Many images of spaces may offer inaccurate portrayals of the relevant properties (size, shape, materials, etc.), or may be misleading (exam-

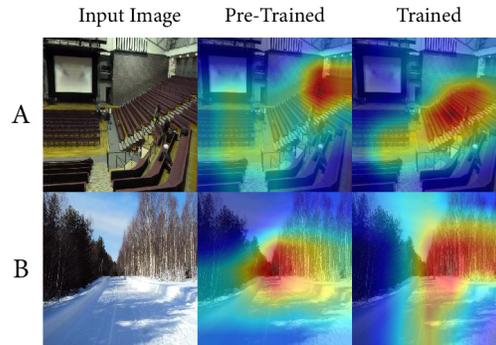


Figure 12. Grad-CAMs for images passed through both the pre-trained Places365 ResNet50 encoder and our fine-tuned encoder, showing movement towards more textured areas for (A) an indoor, and (B) an outdoor environment. The former seems to contain significant absorption and the latter has few reflective surfaces. In both cases, textured areas are highlighted. These may be associated with absorption, diffusion, and more sparse reflections depending on the scene.

ples in supplementary material), leading to erroneous estimations. Our dataset also contains much variation in other relevant parameters (e.g. *DRR* and *EDT*) in a way we cannot semantically connect to paired images, given the sources of our data. New audio IR datasets collected with strongly corresponding photos may allow us to effectively model these characteristics precisely.

5. Conclusion

We introduced Image2Reverb, a system that is able to directly synthesize audio impulse responses from single images. These are directly applied in downstream convolution reverb settings to simulate depicted environments, with applications to XR, music production, television and film post-production, video games, videoconferencing, and other media. Our quantitative and human-expert evaluation shows significant strengths, and we discuss the method’s limitations. We demonstrate that end-to-end image-based synthesis of plausible audio impulse responses is feasible, given such diverse applications. We hope our results provide a helpful benchmark for the community and future work and inspire creative applications.

Acknowledgements We thank the reviewers for their thorough feedback and useful suggestions. Additionally, we thank James Traer, Dor Verbin, and Phillip Isola for helpful discussions. We thank Google for a cloud platform education grant.

References

- [1] Rolf Andereg, Norbert Felber, Wolfgang Fichtner, and Ulrich Franke. Implementation of high-order convolution algorithms with low latency on silicon chips. In *Audio Engi-*

- neering Society Convention 117. Audio Engineering Society, 2004.
- [2] Roland Badeau. Common mathematical framework for stochastic reverberation models. *The Journal of the Acoustical Society of America*, 145(4):2733–2745, 2019.
 - [3] Nicholas J Bryan. Impulse response data augmentation and deep neural networks for blind room acoustic parameter estimation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2020.
 - [4] Lele Chen, Sudhanshu Srivastava, Zhiyao Duan, and Chenliang Xu. Deep cross-modal audio-visual generation. *CoRR*, abs/1704.08292, 2017.
 - [5] Terrance DeVries, Adriana Romero, Luis Pineda, Graham W. Taylor, and Michal Drozdal. On the evaluation of conditional {gan}s, 2020.
 - [6] Chris Donahue, Julian McAuley, and Miller Puckette. Adversarial audio synthesis. In *International Conference on Learning Representations*, 2019.
 - [7] Jesse Engel, Kumar Krishna Agrawal, Shuo Chen, Ishaan Gulrajani, Chris Donahue, and Adam Roberts. GANSynth: Adversarial neural audio synthesis. In *International Conference on Learning Representations*, 2019.
 - [8] Ruohan Gao and Kristen Grauman. 2.5d visual sound. In *CVPR*, 2019.
 - [9] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth prediction. In *The International Conference on Computer Vision (ICCV)*, October 2019.
 - [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
 - [11] Samuel W Greenhouse and Seymour Geisser. On methods in the analysis of profile data. *Psychometrika*, 24(2):95–112, 1959.
 - [12] Jie Gui, Zhenan Sun, Yonggang Wen, Dacheng Tao, and Jieping Ye. A review on generative adversarial networks: Algorithms, theory, and applications, 2020.
 - [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
 - [14] H. Kim, L. Remaggi, P. J. B. Jackson, and A. Hilton. Immersive spatial audio reproduction for vr/ar using room acoustic modelling from 360° images. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 120–126, 2019.
 - [15] Homare Kon and Hideki Koike. Deep neural networks for cross-modal estimations of acoustic reverberation characteristics from two-dimensional images. In *Audio Engineering Society*, 05 2018.
 - [16] Homare Kon and Hideki Koike. Estimation of late reverberation characteristics from a single two-dimensional environmental image using convolutional neural networks. *Journal of the Audio Engineering Society*, 67:540–548, 08 2019.
 - [17] Homare Kon and Hideki Koike. An auditory scaling method for reverb synthesis from a single two-dimensional image. *Acoustical Science and Technology*, 41(4):675–685, 2020.
 - [18] Keun Sup Lee, Nicholas J Bryan, and Jonathan S Abel. Approximating measured reverberation using a hybrid fixed/switched convolution structure. In *Proceedings of the 13th International Conference on Digital Audio Effects (DAFx’10)*, 2010.
 - [19] Dingzeyu Li, Timothy R. Langlois, and Changxi Zheng. Scene-aware audio for 360° videos. *ACM Trans. Graph.*, 37(4), July 2018.
 - [20] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802, 2017.
 - [21] Z. Meng, F. Zhao, and M. He. The just noticeable difference of noise length and reverberation perception. In *2006 International Symposium on Communications and Information Technologies*, pages 418–421, 2006.
 - [22] Fabian Mentzer, George D Toderici, Michael Tschannen, and Eirikur Agustsson. High-fidelity generative image compression. *Advances in Neural Information Processing Systems*, 33, 2020.
 - [23] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
 - [24] Damian T Murphy and Simon Shelley. Openair: An interactive auralization web resource and database. In *Audio Engineering Society Convention 129*. Audio Engineering Society, 2010.
 - [25] Egon S Pearson, Ralph B D “AGOSTINO, and Kimiko O Bowman. Tests for departure from normality: Comparison of powers. *Biometrika*, 64(2):231–246, 1977.
 - [26] Renato S Pellegrini. Quality assessment of auditory virtual environments. In *Proceedings of the 2001 International Conference on Auditory Display*, 2001.
 - [27] Anton Ratnarajah, Zhenyu Tang, and Dinesh Manocha. Irgan: Room impulse response generator for speech augmentation, 2021.
 - [28] Andrew Reilly and David McGrath. Convolution processing for realistic reverberation. In *Audio Engineering Society Convention 98*. Audio Engineering Society, 1995.
 - [29] Luca Remaggi, Hansung Kim, Philip JB Jackson, and Adrian Hilton. Reproducing real world acoustics in virtual reality using spherical cameras. In *Proceedings of the 2019 AES International Conference on Immersive and Interactive Audio*. Audio Engineering Society, 2019.
 - [30] Michael Rettinger. Reverberation chambers for broadcasting and recording studios. *Journal of the Audio Engineering Society*, 5(1):18–22, 1957.
 - [31] H Robjohns. Sony dre s777 sampling digital reverb. *Sound on Sound*, 15, 1999.
 - [32] Kaushal Sali and Alexander Lerch. Generating impulse responses using recurrent neural networks, 2020.
 - [33] Carl Schissler and Dinesh Manocha. Interactive sound propagation and rendering for large multi-source scenes. *ACM Trans. Graph.*, 36(4), Sept. 2016.

- [34] M. R. Schroeder and B. F. Logan. “Colorless” Artificial Reverberation. *IRE Transactions on Audio*, 9(6):209–214, 1961.
- [35] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [36] Christian Steinmetz. Neuralreverberator, 2018.
- [37] James Traer and Josh H McDermott. Statistics of natural reverberation enable perceptual separation of sound and space. *Proceedings of the National Academy of Sciences*, 113(48):E7856–E7865, 2016.
- [38] Vesa Valimaki, Julian D Parker, Lauri Savioja, Julius O Smith, and Jonathan S Abel. Fifty years of artificial reverberation. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(5):1421–1448, 2012.
- [39] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alexander Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. In *Arxiv*, 2016.
- [40] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [41] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. *Advances in Neural Information Processing Systems*, 1(January):487–495, 2014.