

Beyond Image to Depth: Improving Depth Prediction using Echoes

Kranti Kumar Parida¹

Siddharth Srivastava²

Gaurav Sharma^{1,3}

¹ IIT Kanpur

² CDAC Noida

³ TensorTour Inc.

{kranti, grv}@cse.iitk.ac.in, siddharthsrivastava@cdac.in

Abstract

We address the problem of estimating depth with multi-modal audio visual data. Inspired by the ability of animals, such as bats and dolphins, to infer distance of objects with echolocation, some recent methods have utilized echoes for depth estimation. We propose an end-to-end deep learning based pipeline utilizing RGB images, binaural echoes and estimated material properties of various objects within a scene. We argue that the relation between image, echoes and depth, for different scene elements, is greatly influenced by the properties of those elements, and a method designed to leverage this information can lead to significantly improved depth estimation from audio visual inputs. We propose a novel multi modal fusion technique, which incorporates the material properties explicitly while combining audio (echoes) and visual modalities to predict the scene depth. We show empirically, with experiments on Replica dataset, that the proposed method obtains 28% improvement in RMSE compared to the state-of-the-art audio-visual depth prediction method. To demonstrate the effectiveness of our method on larger dataset, we report competitive performance on Matterport3D, proposing to use it as a multi-modal depth prediction benchmark with echoes for the first time. We also analyse the proposed method with exhaustive ablation experiments and qualitative results. The code and models are available at <https://krantiparida.github.io/projects/bimgdepth.html>

1. Introduction

Humans perceive the surroundings using multiple sensory inputs such as sound, sight, smell and touch, with different tasks involving different combinations of such inputs.

In computer vision, multimodal learning has also gained interest. As one popular stream, researchers have leveraged audio and visual inputs for addressing challenging problems. These problems can be broadly divided into three categories: (i) using audio modality only as the input, to learn a seemingly visual task, e.g. using echo for depth prediction [9], (ii) using visual modality as *auxilliary information* for an audio task, e.g. using videos to convert mono audio to

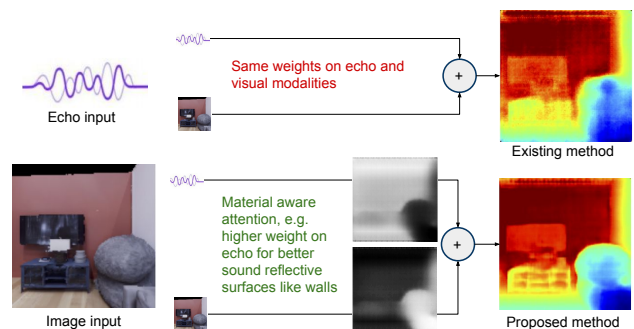


Figure 1. We address the problem of depth prediction using multimodal audio (binaural echo) and visual (monocular RGB) inputs. We propose an attention based fusion mechanisms, where the attention maps are influenced by automatically estimated material properties of the scene objects. We argue that capturing the material properties while fusing echo with images is beneficial as the light and sound reflection characteristics depend not only on the depth, but also on the material of the scene elements.

binaural audio [14], and (iii) using both audio visual modalities together, e.g. for depth prediction [12]. Here, we follow the third line of work, and address the problem of depth map prediction using both audio and visual inputs. Studies in psychology and perception indicate that both sound and vision complement each other, i.e. visual information helps calibrate the auditory information [22] while auditory grouping helps solve visual ambiguity [42]. Many animals, like bats and dolphins, use echolocation to estimate the distances of objects from them. Visually impaired humans have also been reported to use echolocation [1]. Motivated by such cases, Christensen et. al [9, 10] recently showed that depth maps can be predicted directly from stereo sound. Gao et. al [12] showed that by fusing features from binaural echoes with the monocular image features, depth estimation can be improved. Inspired by these findings, we work with similar reasoning, i.e. sound contains useful information to predict depth, and that echoes, used along with monocular images, improve depth estimation.

Going beyond the current methods which do simple combinations of features from echoes and images [12], we

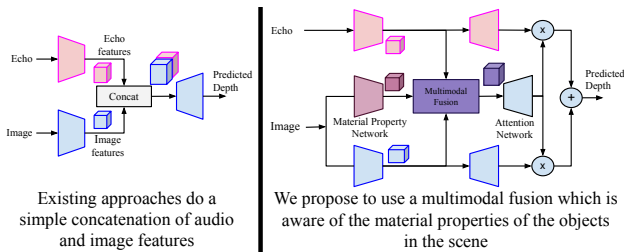


Figure 2. Comparison of our method with the existing approaches

argue that the material properties of the objects in the scene significantly inform the spatial fidelity of the two streams. Some objects may lend better depth estimates with echoes, while some may prefer the visual modality more. Deriving from this motivation, we propose a novel end-to-end learnable network with a multimodal fusion module. This novel module incorporates material properties of the scene and fuses the two modalities with spatial attention maps indicating the fidelity of the respective modality for different spatial locations. The material properties are automatically estimated using a sub-network initialized with training on auxiliary data on materials. As the final depth prediction, the method fuses the depth maps produced by the audio and visual inputs, modulated by the predicted attention maps. Fig. 1 illustrates the difference with a real output of an existing method and the proposed approach, showing qualitative improvements.

We demonstrate the advantages of the proposed method with experiments on Replica [37] and Matterport3D [7] datasets. We outperform the previous state-of-the-art on Replica dataset by $\sim 28\%$ RMSE. On Matterport3D, which is more complex and larger (5x) than Replica, we provide results on the multimodal depth prediction task for the first time, and compare the proposed method with existing approaches and challenging baselines. We also show that the proposed network can estimate better depth with low resolution images. This is important in practical systems working on depth estimation from monocular images, as sensors capturing echoes can be used along with cameras, to not only enhance the performance of existing setup but also suffer lesser degradation in depth prediction with the reduction in the quality of images. Further, we give ablation experiments to systematically evaluate the different aspects of the proposed method.

In summary, we make the following contributions:

- We propose a novel end-to-end learnable deep neural network to estimate depth from binaural audio and monocular images.
- We provide exhaustive quantitative and qualitative results on Replica and Matterport3D datasets. On Replica, we outperform the previous state-of-the-art by $\sim 28\%$. On Matterport3D, we provide results benchmarking existing methods. The proposed method achieves state-

of-the-art performance, outperforming the existing best method on Matterport3D by $\sim 4\%$.

- We provide exhaustive ablation experiments on the design choices in the network, and validate our intuitions with representative qualitative results.

2. Related Works

Audio-visual learning. Recently there has been a surge in interest in audio-visual learning. In one line of work, the correspondence between both the modalities are used to learn the representation in each of the individual modality, in a self-supervised manner [2, 3, 19, 27, 29]. In [2, 3], the authors used an auxiliary task, of predicting whether the audio and image pair correspond to each other, to learn representations in each of the modalities. In [29], the authors predicted if the audio and the video clip are temporally synchronized to learn representations. In [19], the authors advanced a step further and localized sound generating objects in the image by leveraging the correspondence between audio and the image. In one of the recent approach, the authors in [27], have used the spatial correspondence between 360° video and audio. In another line of work, the integration of both audio and video modality was done to increase the performance. Recently a variety of task such audio source separation [49, 48, 11, 13, 15], zero-shot learning [30, 26], saliency prediction [40], audio spatialization [14, 28] have used the information from both audio and video modalities to improve the performance cf. using single modality only.

Depth estimation without echoes. The depth estimation methods span from only monocular image based methods to multi modal methods. Usually, the modalities are sparse depth maps, LiDAR point clouds, bird’s eye views, and normal maps. Monocular depth estimation methods involve utilizing a single RGB image to estimate dense depth [47, 6, 39]. Many methods directly utilize single image [16, 23, 32] or estimate an intermediate 3D representations such as point clouds [43, 45] and bird’s eye views [35] to estimate dense depth maps. A few other methods work on combining RGB with sparse depth maps, normal maps etc. [31, 25] to estimate dense depth maps.

Depth estimation with echoes. In [9, 10] depth of the scene was estimated using only echoes received from a single audio pulse. This approach completely ignored the visual modality while estimating the depth. On similar lines, the authors in [41] estimated the depth of scene directly from binaural audio of the object itself. They did not have the ground truth depth map, and instead used a vision network to predict the ground truth depth map. This method although used the direct audio from the object but the performance of the system was always upper bounded by the predicted depth map from visual input. In all of the methods above, the authors used one modality in isolation and did not fuse multi-modal information to improve the perfor-

mance of the system. In [12] the authors used echolocation as a pre-training task for learning a better visual representation. The authors also gave a case study, where they showed that simply concatenating the audio features with the visual features improves depth prediction. We explore the idea further, of improving depth prediction by adding binaural echoes to image input as well, and propose a novel multimodal fusion method which incorporates automatically estimated material properties in addition. We give the comparison with existing approach in Fig. 2.

3. Approach

We are interested in using echoes for predicting depth map of a scene. There is a spatio-temporal relationship between the received echoes and the depth of the scene, i.e. the echoes received at different instant of time directly relate to the different depths in the scene. Let $x(t)$ be the original audio signal and $y(t)$ be the echo response of the scene. Assuming we have d_i distinct depths of materials m_i in the scene (discretized and grouped for simplicity of discussion), the obtained echo response can be approximated with a summation of time delayed original signal, reflecting off the materials at different depths. The amplitudes of the delayed signals will depend upon the material they hit respectively. Considering only first order echoes, say every distinct d_i object contributes to a time delay of t_i , and the corresponding material changes the amplitude of the signal by a factor of a_i on an average. The final response¹ could then be approximated as

$$y(t) = x(t) + \sum_{i=1}^k a_i x(t - t_i). \quad (1)$$

With v_s denoting the speed of sound in the medium, the time delay can be directly associated with depths as $t_i = \frac{2d_i}{v_s}$. Further, the amplitudes, a_i of each time-delayed signal, are dependent on the acoustic absorption and reflection coefficient of the material.

Hence, the goal of making the network learn depth from the received echo, is influenced by two factors, (i) the relationship between the echoes and spatial depth variation in the scene, and (ii) the different acoustic properties of the scene objects. We propose a carefully tailored attention mechanism (Sec. 3.6) between the image and audio modalities for addressing the spatial variation aspect. In addition we also propose to incorporate material property estimation (Sec. 3.4) as a proxy, to account for different properties of the scene elements which inform sound and light reflection and absorption, and hence inform the final depth prediction.

¹The final response will include the original signal as well, as the sound emitter and recorder are both turned on together for a brief period of time

3.1. Overall architecture

We show the block diagram of the proposed network in Fig. 3. The network consists of the following components: (i) echo subnetwork, (ii) visual subnetwork, (iii) material properties subnetwork, (iv) the multimodal fusion module and finally (v) attention prediction subnetwork. The echo and visual subnetworks consist of encoder-decoder pairs which estimate depth maps of the scenes independently. We input three feature maps coming from echo, visual and material property subnetworks respectively to the multimodal fusion network. The multimodal fusion module produces the fused features which we then feed to the attention network. We further use the attention network to predict two attention maps, one each for the two depth maps obtained from echo and visual decoder networks respectively. We then combine the individual depth maps using these attention maps to output the final depth map of the scene. We now give the details of the different components below. Please also see the supplementary document for detailed layer-wise architecture of the method.

3.2. Echo Net for Echo to Depth

The echo net is an encoder decoder network which predicts depth from binaural echo input. We convert the time-domain echo response into a frequency domain spectrogram representation, $\mathbf{E} \in \mathbb{R}^{2 \times P \times Q}$, where P is the number of discrete time steps and Q is the number of frequency bins. We input the spectrogram to the encoder part of the network to obtain the encoded representation $f_e \in \mathbb{R}^N$ of the echo, which is also one of the inputs to the multimodal fusion module. We then reshape the encoded vector to $N \times 1 \times 1$ and feed it to a series of fractionally strided convolution layers to get the depth map, $\mathbf{D}_e \in \mathbb{R}^{W \times H}$, where W, H are the width and the height of the input image. While the upsampling happens here from an extreme 1×1 feature, i.e. there is practically no spatial information (except what might get coded in the feature vector itself), such depth prediction from audio has been reported by earlier works also with fair success [12].

3.3. Visual Net for Image to Depth

The visual net is also an encoder decoder network which predicts depth from monocular RGB image. The architecture of this network is inspired from U-Net [33], and consists of regular convolutional and fractionally strided convolutional layers with skip connections between them. We give the image, $\mathbf{I} \in \mathbb{R}^{3 \times W \times H}$ as input to the network which predicts the depth map, $\mathbf{D}_i \in \mathbb{R}^{W \times H}$. We also use it to obtain the visual features from the intermediate layer (output of last conv layer) of the network denoted, $f_i \in \mathbb{R}^{N \times w \times h}$. We use this feature as one of the inputs to the multimodal fusion module as well.

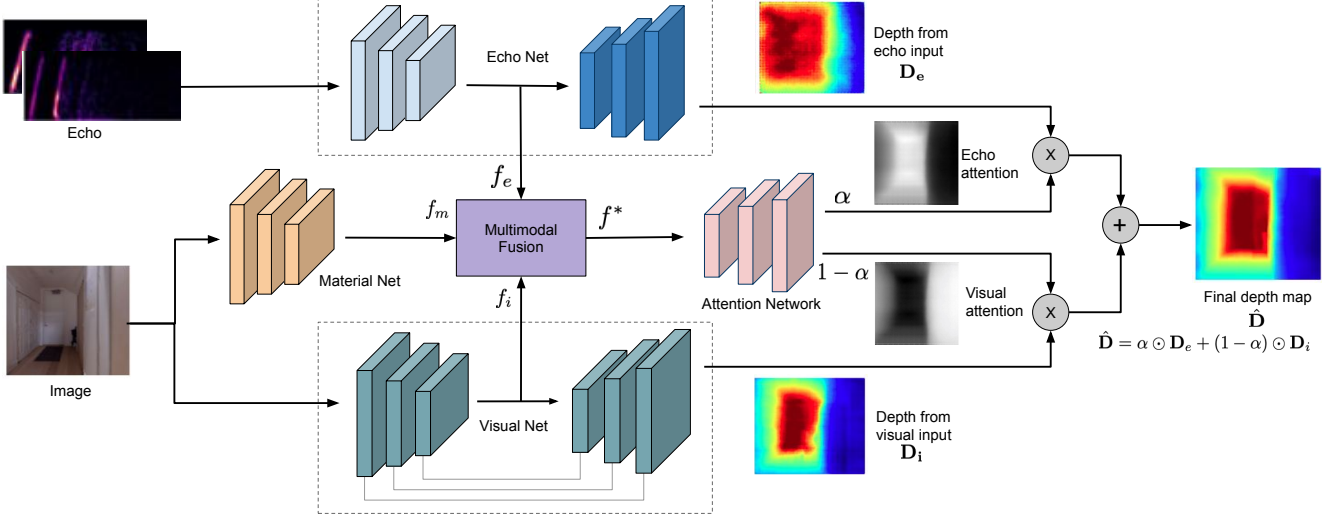


Figure 3. **Block diagram showing the overall architecture of the proposed method.** The network predicts depth independently from binaural echoes and image using Echo Net and Visual Net respectively. The Material Net gives the material features and the multimodal fusion block takes the image feature (f_i), material feature (f_m) and echo feature (f_e) as input and fuses them into a combined feature (f^*). The combined feature (f^*) is then input to attention network to get point-wise attention maps for the echo and image based depths, respectively. The final depth is predicted by taking a weighted combination of echo and image depth with attention maps as the weights.

3.4. Material Net for Material Properties

We use this network to extract the material properties of the objects present in the scene. We use a ResNet-18 architecture [18] and feed the RGB image, $\mathbf{I} \in \mathbb{R}^{3 \times W \times H}$ as the input. We obtain a feature map, $f_m \in \mathbb{R}^{N \times w \times h}$ which encodes the material properties over the spatial locations in the scene image. This feature is the third input to the multimodal fusion module. We initialize the material network with pretraining on a large materials dataset [4] with classes such as fabric, brick, asphalt, wood, metal, and then train it end to end with the rest of the network. We expect this initial material encoding capability of the network to be a proxy for encoding properties related to sound and light absorption and reflection, which affect depth prediction. Although the network evolves with the end to end training, the attention maps obtained qualitatively validate our assumptions (Sec. 4.5).

3.5. Multimodal Fusion Module

The multimodal fusion module combines features from the three sources discussed above, i.e. echo $f_e \in \mathbb{R}^N$, visual $f_i \in \mathbb{R}^{N \times w \times h}$ and material $f_m \in \mathbb{R}^{N \times w \times h}$. Given the motivation, discussed in Sec. 1, that different object might give different depth prediction performances with audio or visual modalities respectively, the multimodal fusion module helps us combine the modalities to provide as input to the attention prediction network (Sec. 3.6).

We perform two bilinear transforms on the features to obtain two fusion maps, f_{img}^j and f_{mat}^j , where $j = 1, 2, \dots, K$ is the number of output channels in the bilinear

transformation,

$$f_{img}^j(p, q) = f_e^T \mathbf{A}_{img}^j f_i(p, q) + b_{img}^j, \forall p, q \quad (2)$$

$$f_{mat}^j(p, q) = f_e^T \mathbf{A}_{mat}^j f_m(p, q) + b_{mat}^j, \forall p, q \quad (3)$$

where, (p, q) indexes the spatial coordinates, \mathbf{A}_{img}^j , \mathbf{A}_{mat}^j are learnable weights of dimension $N \times N$ and b_{img}^j , b_{mat}^j are scalar biases.

We finally concatenate the fusion maps $f_{img} \in \mathbb{R}^{N \times w \times h}$ and $f_{mat} \in \mathbb{R}^{N \times w \times h}$ along the first dimension to get the final fusion map $f^* = \text{concat}(f_{img}, f_{mat})$ to be fed into the per-pixel attention network.

3.6. Attention Network

The attention network is the final part of the network which we use to predict the per-pixel attention map given the concatenated fusion maps obtained in the previous step. The network consists of a series of fractionally strided convolutional layers with a final Sigmoid layer to normalize the values in the range $[0, 1]$.

The output of the network is an attention map $\alpha \in \mathbb{R}^{1 \times W \times H}$. We use the attention map α for weighting the echo predicted depth map D_e and $1 - \alpha$ for the image predicted depth map D_i . The final depth map \hat{D} is thus,

$$\hat{D} = \alpha \odot D_e + (1 - \alpha) \odot D_i \quad (4)$$

where, \odot denotes pointwise multiplication.

3.7. Loss Function and Training

We train the network following [20], and use the logarithm of depth errors. The loss is given as,

$$\mathcal{L}(\hat{\mathbf{D}}, \mathbf{D}) = \frac{1}{WH} \sum_{p=1}^W \sum_{q=1}^H \ln(1 + \|\mathbf{D}(p, q) - \hat{\mathbf{D}}(p, q)\|_1), \quad (5)$$

where D is the ground truth depth map.

The full optimization problem is given by

$$\theta_e^*, \theta_i^*, \theta_a^*, \theta_f^*, \theta_m^* = \arg \min_{\theta_e, \theta_i, \theta_a, \theta_f, \theta_m} \mathcal{L}(\hat{\mathbf{D}}, \mathbf{D}). \quad (6)$$

where, $\theta_e, \theta_i, \theta_a, \theta_f, \theta_m$ are the parameters for echo to depth network, image to depth network, attention network, fusion module and material property network respectively. We ignore the undefined regions in the ground truth depth maps, and therefore, such regions do not contribute to the learning. Adding smoothness constraints [23] can potentially further improve the quality of generated depth, however we obtain good results without using them here. We train the full network in an end to end fashion using standard backpropagation for neural networks.

4. Experiments

4.1. Implementation Details

Datasets. We perform experiments on Replica [37] and Matterport3D [7] datasets. Both the datasets contain indoor scenes. Replica has a total of 18 scenes covering hotels, apartments, rooms and offices. Matterport3D contains 90 scenes. On Replica, we follow [12] and use 15 scenes for training and 3 for testing. On Matterport3D, we use 77 scenes for evaluation, out of which 59, 10 and 8 scenes are used as train, validation and test respectively. We simulate echoes on these datasets by using the precomputed room impulse response (RIR) provided by [8] using 3D simulator Habitat [34] which takes into consideration the scene geometry and the materials present in the scene. We obtain the echoes by convolving input audio signal with the RIR. We use the material recognition dataset MINC [4] for pre-training the material net.

Network Architecture. We use the same architecture for echo encoder, and image to depth encoder and decoder (Visual Net) as [12], for fair comparison and demonstrating the effectiveness of the proposed material and attention networks. We use the first four convolutional layers of ResNet-18 for the material property network. We initialize them with pretraining on ImageNet and MINC dataset.

Input Representation. The input to Visual Net and Material Net is a 128×128 RGB image. For input to Echo Net, we use the spectrogram of 60ms echo signal. For training on Replica, we use a sampling frequency of 44.1 kHz and for

Matterport3D, we use a sampling frequency of 16 kHz. We use Hanning window of length 64 and 32 to compute spectrogram for Replica and Matterport3D respectively. We use FFT size of 512 for both the cases.

Metrics. Following earlier works in depth estimation, we report results on root mean squared error (RMSE), mean relative error (REL), mean \log_{10} error, and the percentage δ_t of pixels with both the relative error and its inverse under threshold t where $t \in \{1.25, 1.25^2, 1.25^3\}$.

Due to space constraints, we provide more details on the datasets, network architecture, parameter settings and evaluation metrics in the supplementary material.

4.2. Experiment Design

We design the experiments below to demonstrate the following points. (i) Using audio and visual modalities together improves performance over either of them. (ii) Using material properties in addition improves further. (ii) Among the different ways to combine the three, i.e. visual, audio and material properties, the proposed attention based fusion performs the best. We demonstrate the first two points with ablation experiments where we combine the inputs by simple concatenation, followed by a decoder to predict the depth map (Sec. 4.3 first part). Further, we demonstrate the third point by comparing combination methods and showing that attention based combination performs the best (Sec. 4.3 second part).

We then compare our full method with existing state of the art approaches (Sec. 4.4). We also show experiments on degrading resolution of image (Sec. 4.6).

4.3. Ablation Study

Combination of echo, image and material properties. We show the results of combining the three inputs with simple concatenation in Tab. 1. With only binaural echo as input, the RMSE is 0.995, which improves to 0.673 when image is added as well. When material property features are used with echo, the RMSE improves to 0.523 i.e. an improvement of $\sim 47\%$ over echo only input and $\sim 22\%$ over image+echo input.

Lastly, when image and material property features are concatenated with echo features (all), the RMSE value further improves to 0.491 i.e. $\sim 50\%$ over echo only input and $\sim 27\%$ over echo+image input.

These experiments validate that even with simple fusion, material properties improve the performance of the system. We attribute the improvement to our intuition that adding material properties explicitly allows the network to internally module the audio and visual features.

In the following we demonstrate the proposed explicit multimodal fusion followed by attention based weighting performs much better than the simple concatenation.

Modality	RMSE (\downarrow)	REL (\downarrow)	log10 (\downarrow)	$\delta_{1.25}$ (\uparrow)	$\delta_{1.25^2}$ (\uparrow)	$\delta_{1.25^3}$ (\uparrow)
echo	0.995	0.638	0.208	0.388	0.599	0.742
echo+img	0.673	0.451	0.146	0.534	0.734	0.845
echo+mat.	0.523	0.282	0.103	0.652	0.839	0.920
all	0.491	0.276	0.098	0.667	0.846	0.924

Table 1. **Depth estimation by combining different modalities** Using echoes only (echo), echoes with image features (echo+img.), echoes with material features (echo+mat.) and combination of echo, image and material features (all). \downarrow indicates lower is better and \uparrow indicates higher is better.

Method	RMSE (\downarrow)	REL (\downarrow)	log10 (\downarrow)	$\delta_{1.25}$ (\uparrow)	$\delta_{1.25^2}$ (\uparrow)	$\delta_{1.25^3}$ (\uparrow)
VisualEchoes [12]	0.346	0.172	0.068	0.798	0.905	0.950
concat	0.259	0.122	0.048	0.867	0.939	0.968
dot	0.262	0.133	0.050	0.853	0.943	0.974
bilinear	0.249	0.118	0.046	0.869	0.943	0.970

Table 2. **Performance of different fusion strategies.** concat refers to the concatenation of all the inputs, dot to fusion by dot product, and bilinear to fusion by bilinear transformation (see Sec. 4.3).

Impact of multimodal fusion and attention. We now validate the efficacy of our audio visual fusion method, which uses a multimodal fusion module to predict attention over the modalities to combine them. We compare the proposed fusion method, denoted `bilinear` with two alternatives, i.e. a simple concatenation of features denoted `concat`, and a dot product based fusion denoted `dot`. All these methods fuse the features and use them to estimate attention weights. We also compare by the fusion method of VisualEchoes [12], which fuses features with concatenation and uses them with a decoder to predict depth map, i.e. it has no attention based fusion.

We show the results in Tab. 2. We observe that `bilinear`, with an RMSE of 0.249, performs best among the compared methods, highlighting that the proposed fusion is better than the simple concatenation or dot product based fusion. We also observe that `concat` performs better than VisualEchoes i.e. 0.259 cf. 0.346 RMSE. This indicates that attention maps (which are present in `concat` but absent in VisualEchoes) are important for better performance.

Fig. 4 further shows the training loss (left) and validation RMSE (right) plots. We observe that VisualEchoes suffers from severe overfitting (much higher val RMSE), which is mitigated on adding the material features (i.e. `concat`). This further reinforces the hypothesis that material properties play an important role in depth prediction.

To conclude, we demonstrated from the ablation experiments that, (i) adding material properties explicitly is helpful for audio visual depth prediction, (ii) the proposed fusion strategy is better than simpler alternatives, and (iii) attention based combination of depth maps is better than simple concatenation as used in previous methods, e.g. VisualEchoes.

4.4. Comparison to state-of-the-art

Baselines. We compare on Replica and Matterport3D against VisualEchoes [12] and competitive baselines. The

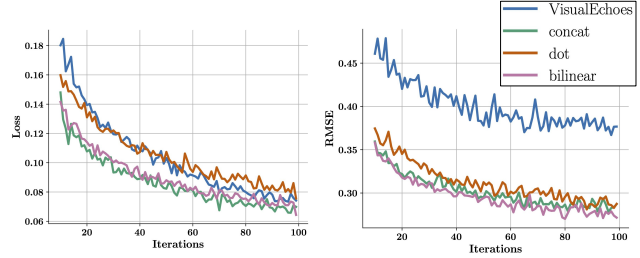


Figure 4. **Training loss and validation RMSE on Replica dataset.** In VisualEchoes [12] depth prediction is performed directly from concatenated features. dot, concat and bilinear are the three different fusion strategies for proposed attention based prediction.

baseline methods are AVERAGE, ECHO2DEPTH and RGB2DEPTH. AVERAGE refers to average depth value of all the samples in training set. ECHO2DEPTH refers to the depth estimation using only Echo Net (Sec. 3.2) and RGB2DEPTH refers to depth estimation using only Visual Net (Sec. 3.3).

Comparison on Replica dataset. We report results in Tab. 3. The proposed method outperforms all the compared methods on all the metrics. Specifically, it outperforms VisualEchoes by $\sim 28\%$ on RMSE. We also observe that while the improvement of VisualEchoes w.r.t. RGB2DEPTH is marginal (0.346 cf. 0.374 i.e. 7.4%), the proposed method is able to achieve an improvement of $\sim 33\%$ (0.249 cf. 0.374 RMSE). Both the methods perform significantly better than AVERAGE and ECHO2DEPTH baselines.

Comparison on Matterport3D dataset. We report results in Tab. 4. We outperform echo only (ECHO2DEPTH), image only (RGB2DEPTH) and AVERAGE baselines on all the five metrics. Our method also outperforms the existing VisualEchoes method by $\sim 4\%$ on RMSE and on all the metric after training the method for 300 epochs. Further, better results on δ indicate that the proposed method has lower pixel wise relative error cf. VisualEchoes which are manifested in form of better depth estimation around edges (Sec. 4.5).

Since Matterport3D is a popular benchmark for depth estimation, we also compare our method with the state-of-the-art methods on sparse to dense depth map estimation. These methods use sparse depth maps as inputs, while we have no explicit depth information in the inputs. We also use a slightly smaller subset of Matterport3D, i.e. 77 cf. 90 scenes for other. The results are shown in Tab. 5 where we obtain better results than four out of five compared methods. While the performances are not directly comparable, it supports the argument that echo can be a viable modality for estimating depth from RGB and can potentially provide additional information that are usually obtained from explicit 3D representations such as sparse depth maps.

Method	RMSE (\downarrow)	REL (\downarrow)	log10 (\downarrow)	$\delta_{1.25}$ (\uparrow)	$\delta_{1.25^2}$ (\uparrow)	$\delta_{1.25^3}$ (\uparrow)
AVERAGE	1.070	0.791	0.230	0.235	0.509	0.750
ECHO2DEPTH	0.713	0.347	0.134	0.580	0.772	0.868
RGB2DEPTH	0.374	0.202	0.076	0.749	0.883	0.945
VisualEchoes [12]	0.346	0.172	0.068	0.798	0.905	0.950
Proposed Method	0.249	0.118	0.046	0.869	0.943	0.970

Table 3. **Comparison with existing methods on Replica dataset.** We report the results for the baseline and existing methods directly from [12].

Method	RMSE (\downarrow)	REL (\downarrow)	log10 (\downarrow)	$\delta_{1.25}$ (\uparrow)	$\delta_{1.25^2}$ (\uparrow)	$\delta_{1.25^3}$ (\uparrow)
AVERAGE	1.913	0.714	0.237	0.264	0.538	0.697
ECHO2DEPTH	1.778	0.507	0.192	0.464	0.642	0.759
RGB2DEPTH	1.090	0.260	0.111	0.592	0.802	0.910
VisualEchoes [12]	0.998	0.193	0.083	0.711	0.878	0.945
Proposed Method	0.950	0.175	0.079	0.733	0.886	0.948

Table 4. **Comparison with existing methods on Matterport3D dataset.**

Method	RMS (\downarrow)	MAE (\downarrow)
AD [24]	1.653	0.610
MRF [17]	1.675	0.618
Zhang et al. [46]	1.316	0.461
Huang et al. [21]	1.092	0.342
Xiong et al. [44]	0.860	0.462
Proposed Method*	1.008	0.570

Table 5. **Comparison on Matterport3D.** *The compared methods use sparse depth maps as inputs, while we do not.




Figure 5. **Visualizing attention map.** We show attention maps for echo and image. First two columns are examples from Replica dataset and last two columns are examples from Matterport3D dataset. We observe that the echo modality, in general, produces high attention value for far away solid structure whereas the image modality attends more to nearby points (sofa in first and third example). See supplementary material for more qualitative results.

4.5. Qualitative Results

We give qualitative results of both (i) depth prediction, and (ii) attention maps in this section. We first provide few qualitative examples for attention map of echo and image respectively in Fig. 5. We provide the attention maps for two examples each from Replica and Matterport3D. We make a general observation from these results that the echo attention map mostly looks into the far off regular regions but completely ignores the finer structures in them. It gives higher weight to the walls in the scenes but ignores any objects present on it (e.g. monitor on third example and wall

Method \ Scale	1x	$\frac{1}{2}$ x	$\frac{1}{4}$ x	$\frac{1}{8}$ x	$\frac{1}{16}$ x	$\frac{1}{32}$ x
Img Only	0.374	0.362	0.375	0.398	0.440	0.550
VisualEchoes[12]	0.346	0.354	0.357	0.392	0.471	0.593
Proposed Method	0.249	0.244	0.249	0.281	0.342	0.446

Table 6. **Performance on varying input image resolution.** The metric is RMSE (lower is better).

painting on fourth). It also ignores the chairs and sofa which are present on relatively lower depth values.

The image attention map, being complementary to the echo attention, gives higher weights to nearby objects such as the sofa, chair etc., but ignores the far off wall in all the examples. This could be mainly due to two reasons (i) the echoes from the nearby regions are observed almost instantaneously which is hard to distinguish from the original signal (ii) most of the nearby objects (e.g. sofa) are made up of sound absorbing materials which do not reflect audio signals strongly. This results suggest that our network tries to leverage the best of both worlds, to get the final depth map.

We give a few examples of reconstructed depth using our approach and also compare it with an existing approach in Fig. 6. We observe that our approach is quite robust to fine grained structure in the scene cf. VisualEchoes [12]. In the first example, VisualEchoes misses the small chairs in the scene while our method gives an idea of the structure. Again the boundaries in the first example are not very clearly identified by VisualEchoes, but are estimated almost perfectly in our case. We also observe similar trends in other three examples as well. The results from the individual modalities (image and echo) are not satisfactory, but do capture the overall structure of the scene. These results suggests that the networks are not individually powerful but their performance improves significantly when we effectively combine the information from both of them. We encourage the readers to look at the supplementary material for more such qualitative results.

4.6. Experiments with varying image resolutions

As discussed in Sec. 4.3, the network can better exploit the combination of echo and image due to attention based integration. This motivates us to study the effect of resolution degradation. This is similar to the case where human vision degrades and the brain learns to compensate and adapt based on the auditory input [5, 38]. We evaluate the proposed method by progressively degrading the image resolution. We give the results in Tab. 6 by gradually reducing the resolution to $\frac{1}{2}$, $\frac{1}{4}$, $\frac{1}{8}$, $\frac{1}{16}$ and $\frac{1}{32}$ times the original image.

We observe that the proposed approach is more robust to reduction in resolution of the images than the compared methods. The performance of image only method degrades significantly when the downscaling factor is $\frac{1}{8}$ of the original image size, while the proposed method still performs better than image only method with original resolution i.e. 0.281 RMSE for proposed method at $\frac{1}{8}$ x scale cf. 0.374 with

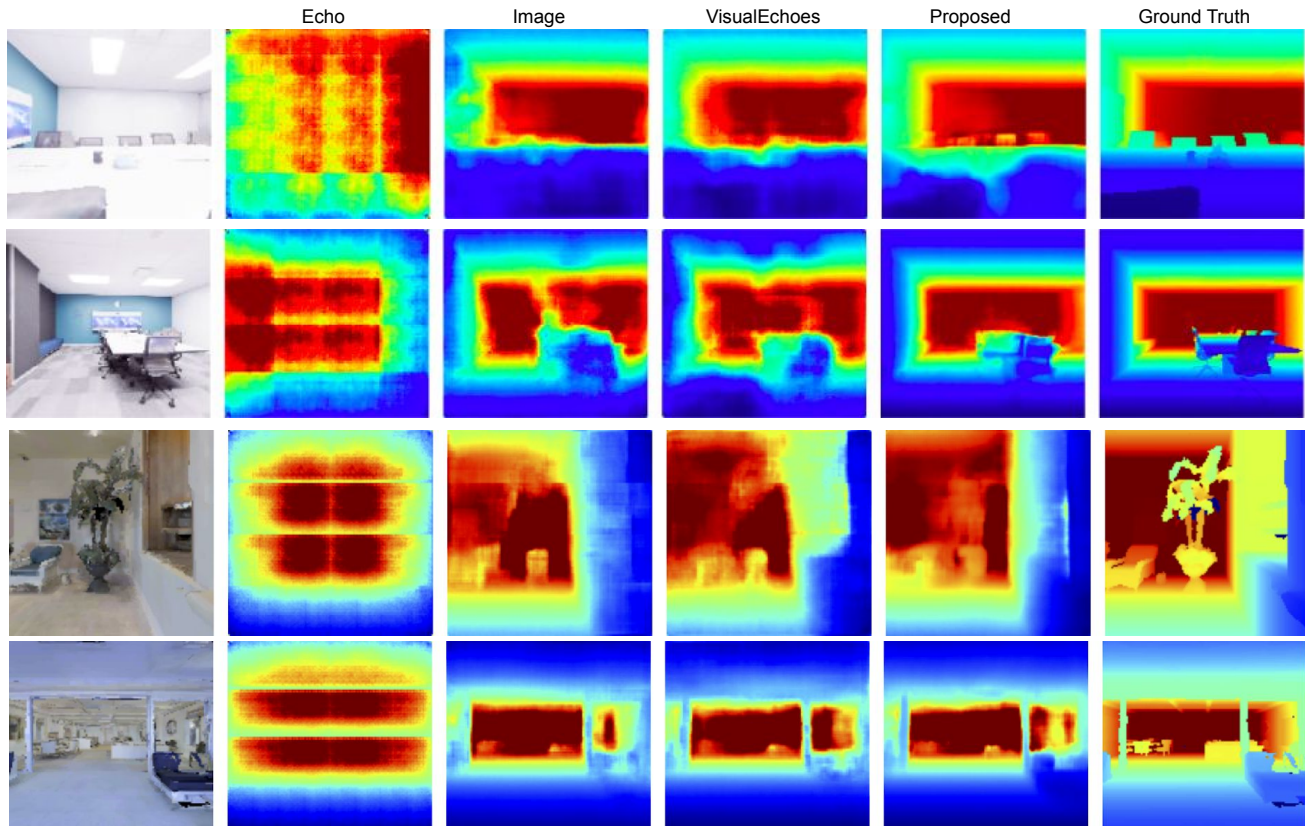


Figure 6. **Qualitative comparisons for depth estimation on Replica (first two rows) and Matterport3D (last two rows) datasets.** We observe that the proposed approach is able to preserve fine structures, and has better depth estimation at the boundaries when compared to the existing approach. See supplementary material for more qualitative results.

image only at $1x$. Further, we observe that even with very high downscaling of $\frac{1}{32}x$, we obtain a better RMSE as compared to VisualEchoes (0.446 cf. 0.593). In fact, VisualEchoes performs worse than even the image only method. Similar observations can be made at $\frac{1}{16}x$. We can also observe that the rate of decrease in RMSE from $1x$ to $\frac{1}{4}x$ is more in VisualEchoes cf. the image only and the proposed method. This further highlights the efficacy of the proposed method.

5. Conclusion

We presented a novel method for estimating depth by combining audio and visual modalities. We hypothesised that material properties play a significant role in the task, and proposed to use automatic material property estimation to predict spatial attention maps which modulate and combine the outputs of audio and image based depth prediction. We showed with quantitative experiments on challenging benchmark datasets, that (i) adding material properties explicitly improves the depth prediction over audio and visual prediction, (ii) having an attention mechanism based fusion method is better than other simpler existing approaches for audio visual fusion. We also demonstrated qualitatively that the attention maps focus on interpretable areas for the two

modalities. While audio attention maps tended to ignore materials which would diffuse or absorb the audio wave, the image based attention included those areas. We also demonstrated qualitatively that the proposed method performs better than existing method, especially near the depth edges and brings out the finer structures in the scene.

We further showed experiments with reduced image resolution where our method degraded gracefully, while the compared methods loses performances significantly. We even compared our method with existing methods for sparse to dense depth prediction, and obtained encouraging competitive results, while not using sparse dense data as input for our method. We would like to explore such multimodal fusion with other modalities like sparse point clouds in the future to obtain even higher quality depth predictions. Further, geometric prior [36] can also be leveraged to improve the results.

In conclusion, we believe that using echo for depth prediction, especially in combination with other modalities is a promising direction, especially given the low cost and wide availability of audio sensors.

Acknowledgment. Kranti Kumar Parida gratefully acknowledges support from the Visvesvaraya fellowship.

References

- [1] https://en.wikipedia.org/wiki/Human_echolocation. 1
- [2] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 609–617, 2017. 2
- [3] Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 2
- [4] Sean Bell, Paul Upchurch, Noah Snavely, and Kavita Bala. Material recognition in the wild with the materials in context database. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3479–3487, 2015. 4, 5
- [5] Jonathan S Berry, David AT Roberts, and Nicolas S Holliman. 3d sound and 3d image interactions: a review of audio-visual depth perception. In *Human Vision and Electronic Imaging XIX*, volume 9014, page 901409. International Society for Optics and Photonics, 2014. 7
- [6] Amlaan Bhoi. Monocular depth estimation: A survey. *arXiv preprint arXiv:1901.09402*, 2019. 2
- [7] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D data in indoor environments. *International Conference on 3D Vision (3DV)*, 2017. 2, 5
- [8] Changan Chen, Unnat Jain, Carl Schissler, Sebastia Vicente, Amengual Gari, Ziad Al-Halah, Vamsi Krishna Ithapu, Philip Robinson, and Kristen Grauman. Soundspaces: Audio-visual navigation in 3d environments. *ECCV*, 2020. 5
- [9] Jesper Haahr Christensen, Sascha Hornauer, and Stella Yu. Batvision: Learning to see 3d spatial layout with two ears. *ICRA*, 2020. 1, 2
- [10] Jesper Haahr Christensen, Sascha Hornauer, and Stella Yu. Batvision with gcc-phat features for better sound to vision predictions. *CVPR Workshop*, 2020. 1, 2
- [11] Chuang Gan, Deng Huang, Hang Zhao, Joshua B Tenenbaum, and Antonio Torralba. Music gesture for visual sound separation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10478–10487, 2020. 2
- [12] Ruohan Gao, Changan Chen, Ziad Al-Halah, Carl Schissler, and Kristen Grauman. Visualechoes: Spatial image representation learning through echolocation. *ECCV*, 2020. 1, 3, 5, 6, 7
- [13] Ruohan Gao, Rogerio Feris, and Kristen Grauman. Learning to separate object sounds by watching unlabeled video. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 35–53, 2018. 2
- [14] Ruohan Gao and Kristen Grauman. 2.5 d visual sound. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 324–333, 2019. 1, 2
- [15] Ruohan Gao and Kristen Grauman. Co-separating sounds of visual objects. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3879–3888, 2019. 2
- [16] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 3828–3838, 2019. 2
- [17] Alastair Harrison and Paul Newman. Image and sparse laser fusion for dense scene reconstruction. In *Field and Service Robotics*, pages 219–228. Springer, 2010. 7
- [18] 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. 4
- [19] Di Hu, Rui Qian, Minyue Jiang, Xiao Tan, Shilei Wen, Errui Ding, Weiyao Lin, and Dejing Dou. Discriminative sounding objects localization via self-supervised audiovisual matching. *Advances in Neural Information Processing Systems*, 33, 2020. 2
- [20] Junjie Hu, Mete Ozay, Yan Zhang, and Takayuki Okatani. Revisiting single image depth estimation: Toward higher resolution maps with accurate object boundaries. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1043–1051. IEEE, 2019. 5
- [21] Yu-Kai Huang, Tsung-Han Wu, Yueh-Cheng Liu, and Winston H Hsu. Indoor depth completion with boundary consistency and self-attention. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. 7
- [22] Andrew J Kolarik, Brian CJ Moore, Pavel Zahorik, Silvia Cirstea, and Shahina Pardhan. Auditory distance perception in humans: a review of cues, development, neuronal bases, and effects of sensory loss. *Attention, Perception, & Psychophysics*, 78(2):373–395, 2016. 1
- [23] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T Freeman. Learning the depths of moving people by watching frozen people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4521–4530, 2019. 2, 5
- [24] Junyi Liu and Xiaojin Gong. Guided depth enhancement via anisotropic diffusion. In *Pacific-Rim conference on multimedia*, pages 408–417. Springer, 2013. 7
- [25] Fangchang Ma, Guilherme Venturelli Cavalheiro, and Sertac Karaman. Self-supervised sparse-to-dense: Self-supervised depth completion from lidar and monocular camera. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3288–3295. IEEE, 2019. 2
- [26] Pratik Mazumder, Pravendra Singh, Kranti Kumar Parida, and Vinay P Namboodiri. Avgzslnet: Audio-visual generalized zero-shot learning by reconstructing label features from multi-modal embeddings. In *The IEEE Winter Conference on Applications of Computer Vision*, 2021. 2
- [27] Pedro Morgado, Yi Li, and Nuno Nvasconcelos. Learning representations from audio-visual spatial alignment. *Advances in Neural Information Processing Systems*, 33, 2020. 2
- [28] Pedro Morgado, Nuno Nvasconcelos, Timothy Langlois, and Oliver Wang. Self-supervised generation of spatial audio for 360 video. In *Advances in Neural Information Processing Systems*, pages 362–372, 2018. 2

- [29] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 631–648, 2018. 2
- [30] Kranti Parida, Neeraj Matiyali, Tanaya Guha, and Gaurav Sharma. Coordinated joint multimodal embeddings for generalized audio-visual zero-shot classification and retrieval of videos. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 3251–3260, 2020. 2
- [31] Jiaxiong Qiu, Zhaopeng Cui, Yinda Zhang, Xingdi Zhang, Shuaicheng Liu, Bing Zeng, and Marc Pollefeys. Deeplidar: Deep surface normal guided depth prediction for outdoor scene from sparse lidar data and single color image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3313–3322, 2019. 2
- [32] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020. 2
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 3
- [34] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9339–9347, 2019. 5
- [35] Siddharth Srivastava, Frederic Jurie, and Gaurav Sharma. Learning 2d to 3d lifting for object detection in 3d for autonomous vehicles. *IROS*, 2019. 2
- [36] Siddharth Srivastava and Gaurav Sharma. Exploiting local geometry for feature and graph construction for better 3d point cloud processing with graph neural networks. In *ICRA*, 2021. 8
- [37] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. 2, 5
- [38] Lore Thaler and Melvyn A Goodale. Echolocation in humans: an overview. *Wiley Interdisciplinary Reviews: Cognitive Science*, 7(6):382–393, 2016. 7
- [39] Lokender Tiwari, Pan Ji, Quoc-Huy Tran, Bingbing Zhuang, Saket Anand, and Manmohan Chandraker. Pseudo rgb-d for self-improving monocular slam and depth prediction. In *European Conference on Computer Vision*, pages 437–455. Springer, 2020. 2
- [40] Antigoni Tsiami, Petros Koutras, and Petros Maragos. Stavis: Spatio-temporal audiovisual saliency network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4766–4776, 2020. 2
- [41] Arun Balajee Vasudevan, Dengxin Dai, and Luc Van Gool. Semantic object prediction and spatial sound super-resolution with binaural sounds. *ECCV*, 2020. 2
- [42] Katsumi Watanabe and Shinsuke Shimojo. When sound affects vision: effects of auditory grouping on visual motion perception. *Psychological Science*, 12(2):109–116, 2001. 1
- [43] Xinshuo Weng and Kris Kitani. Monocular 3d object detection with pseudo-lidar point cloud. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. 2
- [44] Xin Xiong, Haipeng Xiong, Ke Xian, Chen Zhao, Zhiguo Cao, and Xin Li. Sparse-to-dense depth completion revisited: Sampling strategy and graph construction. In *European Conference on Computer Vision (ECCV)*, volume 3, 2020. 7
- [45] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving. *arXiv preprint arXiv:1906.06310*, 2019. 2
- [46] Yinda Zhang and Thomas Funkhouser. Deep depth completion of a single rgb-d image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 175–185, 2018. 7
- [47] Chaoqiang Zhao, Qiyu Sun, Chongzhen Zhang, Yang Tang, and Feng Qian. Monocular depth estimation based on deep learning: An overview. *Science China Technological Sciences*, pages 1–16, 2020. 2
- [48] Hang Zhao, Chuang Gan, Wei-Chiu Ma, and Antonio Torralba. The sound of motions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1735–1744, 2019. 2
- [49] Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba. The sound of pixels. In *Proceedings of the European conference on computer vision (ECCV)*, pages 570–586, 2018. 2