The Boombox: Visual Reconstruction from Acoustic Vibrations

Boyuan Chen

Mia Chiquier Hod Lipson Columbia University boombox.cs.columbia.edu Carl Vondrick

Abstract

We introduce The Boombox, a container that uses acoustic vibrations to reconstruct an image of its inside contents. When an object interacts with the container, they produce small acoustic vibrations. The exact vibration characteristics depend on the physical properties of the box and the object. We demonstrate how to use this incidental signal in order to predict visual structure. After learning, our approach remains effective even when a camera cannot view inside the box. Although we use low-cost and low-power contact microphones to detect the vibrations, our results show that learning from multi-modal data enables us to transform cheap acoustic sensors into rich visual sensors. Due to the ubiquity of containers, we believe integrating perception capabilities into them will enable new applications in human-computer interaction and robotics.

1. Introduction

Reconstructing the occluded contents of containers is a fundamental computer vision task that underlies a number of applications in assistive technology and robotics [7, 38]. However, despite their ubiquity in natural scenes and the ease at which people understand containment [16, 1], containers have remained a key challenge in machine perception [13, 9]. For any camera based task, once an object is contained, there are very few visual signals to reveal the location and appearance of occluded objects.

Recently, the computer vision field has explored several alternative modalities for learning to reconstruct objects occluded by containment. For example, non-line-ofsight imaging systems use the reflections of a laser to sense around corners [6, 11], and radio frequency based sensing shows strong results at visual reconstruction behind walls and other obstructions [41]. These methods leverage the ability to actively emit light or radio frequencies that reflect off surfaces of interest and return to the receiver. These approaches typically require active emission for accurate visual reconstruction.

In this paper, we demonstrate how to use another modal-



Figure 1. **The Boombox:** We introduce a "smart" container that is able to reconstruct an image of its inside contents. Our approach works even when the camera cannot see into the container. The box is equipped with four contact microphones on each face. When objects interact with the box, they cause incidental acoustic vibrations. From just these vibrations, we learn to predict the visual scene inside the box.

ity for reconstructing the visual structure inside containers. Whenever an object or person interacts with a container, they will create an acoustic vibration. The exact incidental vibration produced will depend on the physical properties of the box and its contained objects, such as their relative position, materials, shape, and force. Unlike an active radio, these vibrations are passively and naturally available.

We introduce The Boombox, a smart container that uses the vibration of itself to reconstruct an image of its contents. The box is no larger than a cubic square foot, and it is able to perform all the rudimentary functions that ordinary containers do. Unlike most containers, however, the box uses contact microphones to detect its own vibration. Capitalizing on the link between acoustic and visual structure, we show that a convolutional network can use these vibrations to predict the visual scene inside the container, even under total occlusion and poor illumination. Figure 1 illustrates our box and one reconstruction from the vibration.

Acoustic signals contain extensive information about the surroundings. Humans, for example, use the difference in time and amplitude between both ears to locate sounds and reconstruct shapes [35]. Theoretical results also suggest that, with some assumptions, the geometry of a scene can be reconstructed from audio [14]. However, there are two



Figure 2. **The Boombox Overview.** (A) The Boombox can sense the object through four contact microphones on each side of a storage container. A top-down RGB-D camera is used to collect the final stabilized scene after the object movements. (B) We drop three wooden objects with different shapes. (C) Input and output data visualizations.

key challenges. Firstly, in our setting, the speed of sound is extremely fast for the distance it will travel. Secondly, these methods often assume stationary sound sources, and when they do not, they require high speeds to detect Doppler effects. At the scale of a small box, standard radiolocation methods are not robust because they are sensitive to slight errors in the estimated signal characteristics.

We find that convolutional networks are able to learn robust features that permit reconstruction that is accurate within centimeters. Our experiments demonstrate that acoustics are pivotal for revealing the visual structure behind containment. Given just four microphones attached to each face of the container, we can learn to create an image that predicts both the position and shape of objects inside the box from vibration alone. Our approach works on real, natural audio. Although we use low-cost and low-power microphones, learning from visual synchronization enables us to transform cheap acoustic sensors into 3D visual sensors.

The main contribution of this paper is an integrated hardware and software platform for using acoustic vibrations to reconstruct the visual structure inside containers. The remainder of this paper will describe The Boombox in detail. In section 2, we first review background on this topic. In section 3, we introduce our perception hardware, and in section 4, we describe our learning model. Finally, in section 5, we quantitatively and qualitatively analyze the performance and capabilities of our approach. We will open-source all hardware designs, software, models, and data.

2. Related Work

Audio Analysis. The primary features in audio that are used for sound localization [29] are time difference of arrival and level (amplitude) difference. Specifically calcu-

lating these exact features is non-trivial, especially in situations where the signal is not broad-band and in motion [24, 45, 2, 5]. Furthermore, these rough approximations can only be used to localize the object, whereas our goal is to not only localize objects, but also predict the 3D structure, which includes shape and orientation of the object as well as the environment. As such, we develop a model that learns the necessary features for reconstruction.

Vision and Sound. In recent years the field has seen a growing interest in using sound and vision conjunctively. There are works that, given vision, enhance sounds [30, 18], fill in missing sounds [42], and generate sounds entirely from video [32, 43]. Further, there have been recent works in integrating vision and sound to improve recognition of environmental properties [3, 21, 8] and object properties, such as geometry and materials [40, 39]. Lastly, there have been works in using audiovisual data for representation learning [33, 4, 28]. [17] investigates vision and sound in a robot setting where they predict which robot actions caused a sound. There has been work for generating a face given a voice [31] and a scene from ambient sound [37]. In contrast, our work uses sound to predict the 3D visual structure inside a container.

Non-line-of-sight Imaging: Due to the importance of sensing through occlusions and containers, the field has investigated other modalities for visual reconstruction. In non-line-of-sight imaging, there has been extensive work in scene reconstruction by relying on a laser to reflect off surfaces and return to the receiver [6, 11]. There are also audio extensions [25] as well as radio-frequency based approaches [41]. However, these approaches use specialized and often expensive hardware for the best results. Our approach only uses commodity hardware costing less than \$15



Figure 3. **Chaotic Trajectories:** We show examples of objects trajectories as they bounce around the box until becoming stable. The trajectories are chaotic and sensitive to their initial conditions, making them difficult to visually predict. The moving sound source and multiple bounces also create potentially interfering vibrations, complicating the real-world audio signal.

at the time of writing. Due to this low-cost, a key advantage of our approach is that it is durable, adaptable and straightforward, which makes it easy for others to build on.

3. The Boombox

In this section, we present the The Boombox and discuss the characteristics of the acoustic signals captured by it.

3.1. Detecting Vibrations

The Boombox, shown in Figure 2A, is a plastic storage container that is $15.5 \text{cm} \times 26 \text{cm} \times 13 \text{cm}$ (width \times length \times height) with an open top. The box is a standard object that one can buy at any local hardware store.

When an object collides with the box, a small acoustic vibration will be produced in both the air and the solid box itself. In order for the box to detect its own vibration, we have attached contact microphones on each wall of the plastic cuboid storage bin. Unlike air microphones, contact microphones are insensitive to the vibrations in the air (which human ears hear as sound). Instead, they detect the vibration of solid objects.

The microphones are attached on the outer side of the walls, resulting in four audio channels. We arrange the microphones roughly at the horizontal center of each wall and close to the bottom. As our approach will not require calibration, the microphone displacements can be approximate. We used TraderPlus piezo contact microphones, which are very affordable (no more than \$5 each).¹

3.2. Vibration Characteristics

When objects collide with the box, the contact microphones will capture the resulting acoustic vibrations. Figure 4 shows an example of the vibration captured from two of the microphones. We aim to recover the visual structure from this signal. As these vibrations are independent of the



Figure 4. **Vibration Characteristics:** We visualize a vibration captured by two microphones in the box. There are several distinctive characteristics that need to be combined over time in order to accurately reconstruct an image with the right position, orientation, and shape of objects.

visual conditions, they allow perception despite occlusion and poor illumination.

There is rich structure in the raw acoustic signal. For example, the human auditory system uses inter-aural timing difference (ITD), which is the time difference of arrival between both ears, and inter-aural level difference (ILD), which is the amplitude level difference between both ears, to locate sound sources [35].

However, in our settings, extracting these characteristics is challenging. In practice, objects will bounce around in the container before arriving at their stable position, as shown in Figure 3. Each bounce will produce another, potentially interfering vibration. In our attempts to analytically use this signal, we found that the third bounce has the best signal for the time difference of arrival, but as can be seen from Figure 4, even on the third bounce the time difference of arrival is unclear in the actual waveform.

There are a multitude of factors that make analytical approaches not robust to our real-world signals. Firstly, we are working with a moving signal, whereas time difference of arrival calculations work best on stationary signals due to the fact that it compares the time taken for a signal to travel from a fixed location. This makes it very difficult to analytically segment the signal into chunks of roughly the same location. Secondly, there are echos that make non-learning based methods difficult to identify phase shifts as the environment is a small container. Finally, the fact that the microphones are close together means that the time difference of arrival is encompassed in few samples, thus making it susceptible to noise.

Instead of hand crafting these features, our model will learn to identify the fraction of the signal that is most robust for final localization. Moreover, our model will learn to identify the useful features from the signals to reconstruct a rich 3D scene that includes the shape, orientation, and position of the contents.

3.3. Multimodal Training Dataset

Our approach will use the visual modality in order to learn the robust characteristic features of the acoustic signal. By simply dropping objects into the box and capturing resulting images and vibrations, we can collect a mul-

¹We found that these microphones gave sufficiently clear signals while being more affordable than available directional microphone arrays. Each microphone was connected to a laptop through audio jack to USB converter. We use GarageBand software to record all four microphones together to synchronize the recordings.



Figure 5. **Point Cloud Reconstruction from Vibration:** We visualize the point clouds produced by our method. These point clouds are predicted only given the acoustic vibration of the box. For each prediction, we show two different camera views.

timodal training dataset. We position an Intel RealSense D435i camera that looks inside the bin to capture both RGB and depth images.²

We use three wooden blocks with different shapes to create our dataset. The blocks have the same color and materials, and we show these objects in Figure 2B. We hold the object above the bin, and freely drop it.³ After dropping, the objects bounce around in the box a few times before settling into a resting position. We record the full process from all the microphones and the top-down camera. Overall, our collection process results diverse falling trajectories across all shapes with a total of 1,575 sequences. Figure 2C shows an overview of the dataset.

We only use the camera to collect data for learning. After learning, our approach will be able to reconstruct the 3D visual scene from the box's vibration alone.

4. Predicting Images from Vibration

In order to create robust features, we will learn them from multi-modal data. We present a convolutional network that translates vibrations into an image.

4.1. Model

We will fit a model that reconstructs the visual contents from the vibrations. Let A_i be a spectrogram of the vibration captured by microphone *i* such that $i \in \{1, 2, 3, 4\}$. Our model will predict the image $\hat{X}_{RGB} = f_{RGB}(A; \theta)$ where *f* is a neural network parameterized by θ . The network will learn to predict the image of a top-down view into the container. We additionally have a corresponding network to produce a depth image $\hat{X}_{depth} = f_{depth}(A; \theta)$. Reconstructing a pixel requires the model to have access to the full spectrogram. However, we also want to take advantage of the spatio-temporal structure of the signal. We thereforefore use a fully convolutional encoder and decoder architecture. The network transforms a spectrogram representation (time \times frequency) into a C dimensional embedding such that the receptive field of every dimension reaches every magnitude in the input and every pixel in the output. Unlike image-to-image translation problems [44, 20, 12], our task requires translation across modalities.

We use a multi-scale decoder network [27, 19, 10]. Specifically, each decoder layer consists of two branches. One branch is a transposed convolutional layer to up-sample the intermediate feature. The other branch passes the input feature first to a convolutional layer and then a transposed convolution so that the output for the second branch matches the size of the first branch. We then concatenate the output from these two branches along the feature dimension as the input feature for the next decoder layer. We perform the same operation for each decoder layer except the last layer where only one transposed convolution layer is needed to predict the final output image.

We use a spectrogram as the representation of audio signals. We apply a Fourier Transform before converting the generated spectrogram to Mel scale. Since we have four microphones, audio clips are concatenated together along a third dimension in addition to the original time and frequency dimension.

4.2. Learning

In practical applications, we often care about the resting position of the object so that we can localize the object. We therefore train the network f to predict the final stable image. For RGB image predictions, we train the network to minimize the expected mean squared error:

$$\mathcal{L}_{\text{RGB}} = \mathbb{E}_{A,X} \left[\| f_{\text{RGB}}(A;\theta) - X_{\text{RGB}} \|_2^2 \right]$$
(1)

In order to reconstruct shape, we also train the network to predict a depth image from the acoustic vibration input. We train the model to minimize the expected L1 distance:

$$\mathcal{L}_{depth} = \mathbb{E}_{A,X} \left[\| f_{depth}(A;\phi) - X_{depth} \|_1 \right]$$
(2)

Since ground truth depth often has outliers and substantial noise, we use an L1 loss [26]. We use stochastic gradient descent to estimate the network parameters θ and ϕ .

After learning, we can obtain predictions for both the RGB image and the depth image from the acoustic vibrations alone. The visual modality is only supervising representations for the audio modality, allowing reconstructions when cameras are not viable, such during occlusions or low illumination.

²The camera is 42cm away from the bottom of the bin to capture clear top-down images.

³As the dynamics depend on the material properties, we wore a powder free disposable glove while holding the object to avoid changing the humidity on the object surface.



Figure 6. **Single-shape Training Results.** We show the performance of each individual model trained with one of the three objects. We report both the mean and the standard error of the mean from three random seeds. Our approach enables robust features to be learned to predict the location and shape of the dropped objects.



Figure 7. **Multi-shape Training Results.** By mixing all the training data together with all the shapes, our model still outperforms all the baseline methods.

4.3. Implementation Details

Our network takes in the input size of $128 \times 128 \times 4$ where the last dimension denotes the number of microphones. The output is a $128 \times 128 \times 3$ RGB image or a $128 \times 128 \times 1$ depth image. We use the same network architecture for both the RGB and depth output representations except the feature dimension in the last layer for different modalities. All network details are listed in the Appendix.

Our networks are configured in PyTorch [34] and PyTorch-Lightning [15]. We optimized all the networks for 500 epochs with Adam [22] optimizer and batch size of 32 on a single NVIDIA RTX 2080 Ti GPU. The learning rates starts from 0.001 and decrease by 50% at epoch 20, 50, and 100.

5. Experiments

In our experiments, we test the capability of The Boombox to reconstruct an image of its contents from audio inputs. We quantitatively evaluate the model performance. We then show qualitative results for our visual reconstructions. Finally, we show visualizations to analyze the learned representations.

5.1. Dataset

We partition the dataset into a training set (80%), a validation set (10%), and a testing set (10%). We train the

neural networks on the training set, and optimize hyperparameters on the validation set. We report results only on the testing set. All of our results are evaluated on three random seeds for training and evaluation. For each random seed, we also vary the splits of the dataset. We report the mean and the standard error of the mean for all outcomes.

5.2. Evaluation Metrics

Direct measurements in the pixel space is not informative because it is not a perceptual metric. We use two evaluation metrics for our final scene reconstruction that focus on the object state.

IoU measures how well the model reconstructs both shape and location. Since the model predicts an image, we subtract the background to convert the predicted image into a segmentation mask. Similarly, we performed the same operation on the ground-truth image. IoU metric then computes intersection over union with the two binary masks.

Localization score evaluates whether the model produces an image with the block in the right spatial position. With the binary masks obtained in the above process, we can fit a bounding box with minimum area around the object region. We denote the distance between the center of the predicted bounding box and the center of the ground-truth bounding box as d, and the length of the diagonal line of ground-truth box as l. We report the fraction of times the predicted location is less than half the diagonal: $\frac{1}{N} \sum_{i=1}^{N} [d_i \le l/2].$

5.3. Baselines

Time Difference of Arrival (TDoA): We compare against an analytical prediction of the location. In signal processing, the standard practice is to localize sound sources by estimating the time difference of arrival across an array of microphones. In our case, the microphones *surround* the sound source. There are several ways to es-



Figure 8. Model Prediction Visualizations with Unknown Shapes From left to right on each column, we visualize the audio input, the predicted scene, and the ground-truth images. Our model can produce accurate predictions for object shape, position and orientation.



Figure 9. **Performance before and after multi-shape training.** Multi-shape training enables shape knowledge transfer to improve the overall performance

timate the time difference of arrival, and we use the Generalized Cross Correlation with Phase Transform (GCC-PHAT), which is the established, textbook approach [23]. Once we have our time difference of arrival estimate, we find the location in the box that would yield a time difference of arrival that is closest to our estimate.

Random Sampling: To evaluate if the learned models simply memorize the training data, we compared our method against a random sampling procedure. This baseline makes a prediction by randomly sampling an image from the training set and using it as the prediction. We repeated this step for all testing examples over 10 random seeds.

Average Bounding Box: The average bounding box baseline aims to measure to what extent the model learns

the dataset bias. Therefore, we extracted object bounding boxes from all the training data through background subtraction and rectangle fitting to obtain the average center location, box sizes and box orientation. This baseline uses the average bounding box as the prediction for all the test samples.

5.4. Reconstruction via Single Shape Training

In this experiment, we train separate models for each shape of the object independently. Figure 6 shows The Boombox is able to reconstruct both the position and orientation of the shapes. The convolutional network obtains the best performance for most shapes on both evaluation metrics.

Our model performs significantly better than the analytical TDoA baseline. Our method outperforms TDoA often by significant margins, suggesting that our learningbased model is learning robust acoustic features for localization. Due to the realistic complexity of the audio signal, the hand-crafted features are hard to reliably estimate. Our model outperforms both the random sampling and average bounding box baseline, indicating that our model learns the natural correspondence between acoustic signals and visual scene rather than memorizing the training data distribution.

These results highlight the relative difficulty at reconstructing different shapes from sound. By comparing the model performance across various shapes, the model trained on cubes achieves the best performance while the model trained on blocks performs slightly worse. The most difficult shape is the stick.

5.5. Reconstruction via Multiple Shape Training

We next analyze how well The Boombox reconstructs its contents when the shape is not known a priori. We train



Figure 10. Visualization of Ablation Studies: We visualize the impact of different ablations on the model. A) By thresholding the spectrograms, we remove the amplitude from the input. B) Since the differences between the signals from the microphones are important for location, we experimented with flipping the microphones only at testing time. The model's predictions show a corresponding flip as well in the predicted images. C) We also experimented with shifting the relative time difference between the microphones, introducing an artificial delay in the microphones only at testing time. A shift in time causes a shift in space in the model's predictions. The corruptions are consistent with a block falling in that location.

a single model with all the object shapes. The training data for each shape are simply combined together so that the training, validation and testing data are naturally wellbalanced with respect to the shapes. This setting is challenging because the model needs to learn audio features for multiple shapes at once.

We show qualitative predictions for both RGB and depth images in Figure 8. Moreover, since we are predicting a depth image, our model is able to produce a 3D point cloud. We visualize several examples from multiple viewpoints in Figure 5. While analytical approaches are able to predict a 3D scalar position, we are able to predict a 3D point cloud.

Figure 7 shows the convolutional networks are able to learn robust features even when shapes are unknown. When the training data combines all shapes, the model should be able to share features between shapes, thus improving performance. To validate this, we compare performance on the multi-shape versus the single-shape models. We use both IoU and the localization model. Figure 9 shows that the performance on the block and stick shapes are improved by a large margin. We notice that the performance of the cube drops due to the confusion between shapes. When the cube confuses with the stick or the block, because of the



Figure 11. **Quantitative performance under the ablation studies.** We experiment with different perturbations to our input data to understand the model decisions.

smaller surface area of these two shapes, the cube performance slightly degrades.

5.6. Ablations

Since the model generalizes to unseen examples, our results suggest that the convolutional network are learning robust acoustic features. To better understand what features the model has learned specifically, we perform several ablation studies in Figure 10 and Figure 11.

Flip microphones. The microphones' layout should matter for our learned model to localize the objects. When we flipped the microphone location, due to the symmetric nature of the hardware setup, the predictions should also be flipped accordingly. To study this, we flipped the corresponding audio input with a pair-wise strategy, shown in Figure 10. Specifically, the audio input of the Mic1 and Mic4 are flipped, and the audio input of the Mic2 and Mic3 are flipped. Our results in Figure 10B shows that our model indeed produces a flipped scene. The performance in Figure 11 nearly drops to zero, suggesting that the model implicitly learned the relative microphone locations to assist its final prediction.

Remove amplitude. The relative amplitude between microphones is another signal that can indicate the position of the sound source with respect to different microphones. We removed the amplitude information by thresholding the spectrograms, shown in Figure 10. We retrained the network due to potential distribution shift. As expected, even though the time and frequency information are preserved, the model performs much worse (Figure 11), suggesting that our model additionally learns to use amplitude for the predictions.

Temporal shift. We are interested to see if our model learns to capture features about the time difference of arrival between microphones. If time difference of arrival information is helpful, when we shift the audio signal temporally, the model prediction should also shift spatially. We experimented with various degrees of temporal shifts on the original spectrograms. For example, shifting 500 samples corresponds to shifting about 0.01s (500 / 44,000). By shifting the Mic1's spectrogram forward and Mic4's spectrogram backward with zero padding to maintain the same amount of time, and preforming similar operation on Mic2 and Mic3 respectively, we should expect that the predicted object position shifts towards the left-up direction. In Figure 10, we can clearly observe this trend as temporal shift increases. Shifting the signal in time decreases the model's performance, demonstrating that the model has picked up on the time difference of arrival.

5.7. Feature Visualization

We finally visualize the latent features in between our encoder and decoder network by projecting them into a plane with t-SNE[36], shown in Figure 12. We colorize the points according to their ground truth position and orientation. The magnitude distance from the center of the image is represented by saturation, and the angle from the horizontal axis is represented by hue. We find that there is often clear clustering of the embeddings by their position and orientation, showing that the model is robustly discriminating the location of the impact from sound alone. Moreover, the gradual



Figure 12. **t-SNE embeddings** on the latent features from the encoder network. The color encoding follows a color wheel denoting the angle and relative position to the center of the container. Our model learns to encode the position and orientation of the objects in its internal representations.

transitions between colors suggest the features are able to smoothly interpolate spatially.

6. Conclusion

We have introduced The Boombox, a low-cost container integrated with a convolutional network that uses acoustic vibrations to reconstruct a 3D point cloud from image and depth. Containers are ubiquitous, and this paper shows that we can equip them with sound perception to localize and reconstruct their contents inside.

Acknowledgements: We thank Philippe Wyder, Dídac Surís, Basile Van Hoorick and Robert Kwiatkowski for helpful feedback. This research is based on work partially supported by NSF NRI Award #1925157, NSF CA-REER Award #2046910, DARPA MTO grant L2M Program HR0011-18-2-0020, and an Amazon Research Award. MC is supported by a CAIT Amazon PhD fellowship. We thank NVidia for GPU donations. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the sponsors.

References

- Andréa Aguiar and Renée Baillargeon. 2.5-month-old infants' reasoning about when objects should and should not be occluded. *Cognitive psychology*, 39(2):116–157, 1999.
- [2] Inkyu An, Myungbae Son, Dinesh Manocha, and Sung-Eui Yoon. Reflection-aware sound source localization. In 2018 IEEE International Conference on Robotics and Automation (ICRA), pages 66–73. IEEE, 2018.
- [3] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In Proceedings of the IEEE International Conference on Computer Vision, pages 609–617, 2017.
- [4] Yusuf Aytar, Carl Vondrick, and Antonio Torralba. Soundnet: Learning sound representations from unlabeled video. arXiv preprint arXiv:1610.09001, 2016.

- [5] Paolo Bestagini, Marco Compagnoni, Fabio Antonacci, Augusto Sarti, and Stefano Tubaro. Tdoa-based acoustic source localization in the space–range reference frame. *Multidimensional Systems and Signal Processing*, 25(2):337–359, 2014.
- [6] Katherine L Bouman, Vickie Ye, Adam B Yedidia, Frédo Durand, Gregory W Wornell, Antonio Torralba, and William T Freeman. Turning corners into cameras: Principles and methods. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2270–2278, 2017.
- [7] Rodney A Brooks. Planning collision-free motions for pickand-place operations. *The International Journal of Robotics Research*, 2(4):19–44, 1983.
- [8] Fanjun Bu and Chien-Ming Huang. Object permanence through audio-visual representations. *arXiv preprint arXiv:2010.09948*, 2020.
- [9] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8748–8757, 2019.
- [10] Boyuan Chen, Carl Vondrick, and Hod Lipson. Visual behavior modelling for robotic theory of mind. *Scientific Reports*, 11(1):1–14, 2021.
- [11] Wenzheng Chen, Simon Daneau, Fahim Mannan, and Felix Heide. Steady-state non-line-of-sight imaging. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 6790–6799, 2019.
- [12] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8789–8797, 2018.
- [13] Achal Dave, Tarasha Khurana, Pavel Tokmakov, Cordelia Schmid, and Deva Ramanan. Tao: A large-scale benchmark for tracking any object. In *European conference on computer vision*, pages 436–454. Springer, 2020.
- [14] Ivan Dokmanić, Reza Parhizkar, Andreas Walther, Yue M Lu, and Martin Vetterli. Acoustic echoes reveal room shape. *Proceedings of the National Academy of Sciences*, 110(30):12186–12191, 2013.
- [15] WA Falcon and .al. Pytorch lightning. *GitHub. Note:* https://github.com/PyTorchLightning/pytorch-lightning, 3, 2019.
- [16] Lisa Feigenson and Susan Carey. Tracking individuals via object-files: evidence from infants' manual search. *Developmental Science*, 6(5):568–584, 2003.
- [17] Dhiraj Gandhi, Abhinav Gupta, and Lerrel Pinto. Swoosh! rattle! thump!-actions that sound. *arXiv preprint arXiv:2007.01851*, 2020.
- [18] Ruohan Gao and Kristen Grauman. 2.5 d visual sound. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 324–333, 2019.
- [19] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *Pro-*100 (1997) 100

ceedings of the IEEE conference on computer vision and pattern recognition, pages 2462–2470, 2017.

- [20] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [21] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5492–5501, 2019.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- [23] Charles Knapp and Glifford Carter. The generalized correlation method for estimation of time delay. *IEEE transactions* on acoustics, speech, and signal processing, 24(4):320–327, 1976.
- [24] Zhiwei Liang, Xudong Ma, and Xianzhong Dai. Robust tracking of moving sound source using multiple model kalman filter. *Applied acoustics*, 69(12):1350–1355, 2008.
- [25] David B Lindell, Gordon Wetzstein, and Vladlen Koltun. Acoustic non-line-of-sight imaging. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 6780–6789, 2019.
- [26] Fangchang Ma and Sertac Karaman. Sparse-to-dense: Depth prediction from sparse depth samples and a single image. In 2018 IEEE International Conference on Robotics and Automation (ICRA), pages 4796–4803. IEEE, 2018.
- [27] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. arXiv preprint arXiv:1511.05440, 2015.
- [28] Carolyn Matl, Yashraj Narang, Dieter Fox, Ruzena Bajcsy, and Fabio Ramos. Stressd: Sim-to-real from sound for stochastic dynamics. arXiv preprint arXiv:2011.03136, 2020.
- [29] John C Middlebrooks. Sound localization. Handbook of clinical neurology, 129:99–116, 2015.
- [30] Arun Asokan Nair, Austin Reiter, Changxi Zheng, and Shree Nayar. Audiovisual zooming: what you see is what you hear. In Proceedings of the 27th ACM International Conference on Multimedia, pages 1107–1118, 2019.
- [31] Tae-Hyun Oh, Tali Dekel, Changil Kim, Inbar Mosseri, William T Freeman, Michael Rubinstein, and Wojciech Matusik. Speech2face: Learning the face behind a voice. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 7539–7548, 2019.
- [32] Andrew Owens, Phillip Isola, Josh McDermott, Antonio Torralba, Edward H Adelson, and William T Freeman. Visually indicated sounds. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2405–2413, 2016.
- [33] Andrew Owens, Jiajun Wu, Josh H McDermott, William T Freeman, and Antonio Torralba. Ambient sound provides supervision for visual learning. In *European conference on computer vision*, pages 801–816. Springer, 2016.

- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [35] John W Strutt. On our perception of sound direction. *Philosophical Magazine*, 13(74):214–32, 1907.
- [36] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [37] Chia-Hung Wan, Shun-Po Chuang, and Hung-Yi Lee. Towards audio to scene image synthesis using generative adversarial network. In ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 496–500. IEEE, 2019.
- [38] Andy Zeng, Shuran Song, Kuan-Ting Yu, Elliott Donlon, Francois R Hogan, Maria Bauza, Daolin Ma, Orion Taylor, Melody Liu, Eudald Romo, et al. Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching. In 2018 IEEE international conference on robotics and automation (ICRA), pages 3750– 3757. IEEE, 2018.
- [39] Zhoutong Zhang, Qiujia Li, Zhengjia Huang, Jiajun Wu, Joshua B Tenenbaum, and William T Freeman. Shape and material from sound. 2017.
- [40] Zhoutong Zhang, Jiajun Wu, Qiujia Li, Zhengjia Huang, James Traer, Josh H McDermott, Joshua B Tenenbaum, and William T Freeman. Generative modeling of audible shapes for object perception. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1251–1260, 2017.
- [41] Mingmin Zhao, Tianhong Li, Mohammad Abu Alsheikh, Yonglong Tian, Hang Zhao, Antonio Torralba, and Dina Katabi. Through-wall human pose estimation using radio signals. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 7356–7365, 2018.
- [42] Hang Zhou, Ziwei Liu, Xudong Xu, Ping Luo, and Xiaogang Wang. Vision-infused deep audio inpainting. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision, pages 283–292, 2019.
- [43] Yipin Zhou, Zhaowen Wang, Chen Fang, Trung Bui, and Tamara L Berg. Visual to sound: Generating natural sound for videos in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3550–3558, 2018.
- [44] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycleconsistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223– 2232, 2017.
- [45] Mengyao Zhu, Huan Yao, Xiukun Wu, Zhihua Lu, Xiaoqiang Zhu, and Qinghua Huang. Gaussian filter for tdoa

based sound source localization in multimedia surveillance. *Multimedia Tools and Applications*, 77(3):3369–3385, 2018.