

1 Supplementary

2 In this supplementary material we provide the following:

- 3 1. A video for qualitative evaluation of our model’s performance (1.1).
- 4 2. Details regarding AVSpeech-Rooms curation (1.2) (referenced in Sec. 4 of main paper)
- 5 3. Details on our ablation study with different metric training objectives (1.3) (referenced in
- 6 Sec. 5 — "Results on AVSpeech-Rooms" of main paper)
- 7 4. A sample survey slide from our human perception study (1)
- 8 5. Model/training details for our RT60 estimator, de-biased, discriminator, and reverberator
- 9 (1.4) (referenced in Sec. 5 — "Implementation Details" of main paper)
- 10 6. A more detailed version of Figure 4. in the main paper with both baseline models (2)
- 11 7. Details on our data augmentation strategy (1.5) (referenced in Sec. 5 — "Baselines" of main
- 12 paper)
- 13 8. A brief discussion of our work’s limitations and broader impact (1.6 1.7)

14 1.1 Supplementary Video

15 Our video contains several illustrative examples generated by LeMARA on both SoundSpaces-Speech
16 and AVSpeech-Rooms. We provide audio generated by the current state-of-the-art (AViTAR) for
17 reference on each example. We recommend wearing headphones for a better listening experience.

18 1.2 AVSpeech-Rooms

19 Acoustic AVSpeech consists of audio clips from YouTube videos along with an RGB image frame
20 selected randomly from the corresponding video clip. To create AVSpeech-Rooms, we design a set of
21 criteria which we use to filter out samples in which the image contains uninformative, non-natural, or
22 misleading acoustic information about the space. We focus on cases in which the room is not visible,
23 a microphone is being used, or a virtual background/screen is present — any of which will disturb the
24 natural room acoustics for the speaker’s voice. We query each sample with our criteria using a Visual
25 Question Answering (VQA) model [6], which we found more reliable than manual annotations we
26 originally obtained on MTurk. 1 contains information about our criteria.

Table 1: Filtering criteria and % of Acoustic AVSpeech samples removed.

Question	Answer	dataset %
Is a microphone or headset visible in the image?	yes	7.2
Is there a whiteboard/blackboard in the background?	yes	3.4
Is the entire background one solid color and material?	yes	23.4
Is there a large projector screen covering most of the background?	yes	2.2
Is part or all of the background virtual?	yes	1.3
Are there multiple screens in the image?	yes	3.5
Is the wider room clearly visible?	no	3.0

27

28 1.3 Ablations

29 Table 2 displays our experiments with different self-supervised training objectives. We report
30 performance on the LibriSpeech evaluation setting. The first three rows correspond to experiments
31 in which we do not utilize the shortcut training strategy (referenced in Sec. 3 — "Training" of main
32 paper). Using SRMR alone (row 1) produces the largest (worst) RTE. Training with the acoustic
33 residue metric instead (row 2) leads to a large improvement in RTE, providing empirical support for
34 our metric as an effective training objective. Using our combined metric and the shortcut training
35 strategy (both described in Sec. 3 — "Training" of our main paper) further improves the performance
36 by a small margin.

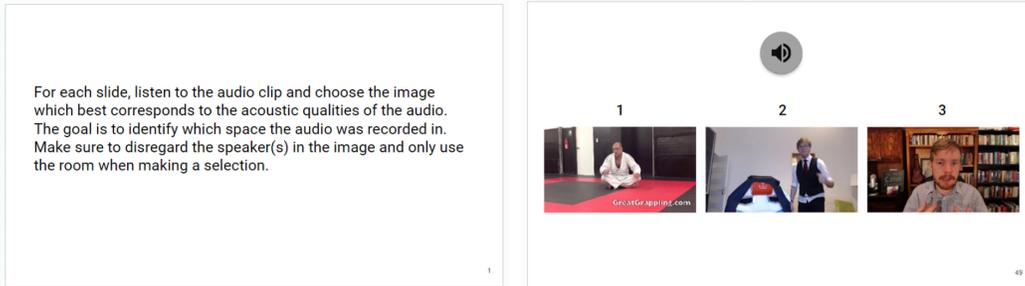


Figure 1: **Human perception study.** The instructions given to the user at the start of the survey (left), and a sample slide from the survey (right). The user is asked to listen to the audio clip, and identify which room image most closely matches its acoustics.

Table 2: Ablation study using different metric training objectives. AR denotes the proposed acoustic residue metric.

Metric	<i>LibriSpeech</i> RTE
SRMR [4]	0.2308
AR	0.2156
AR (combined)	0.2123
AR (combined) w/ shortcut	0.2100

37 1.4 Model/Training details

38 **RT60 estimator** We adopt the RT60 estimator from [1]. The estimator takes a spectrogram as
 39 input, encodes it with a ResNet18 [5], and outputs a scalar RT60 estimate. The model is trained on
 40 2.56s clips of reverberant speech simulated on the SoundSpaces platform [2] paired with the ground
 41 truth RT60 computed from the RIR used to generate the reverberant speech. The model trains using
 42 MSE loss between predicted and ground truth RT60 values. Ground truth RT60 is computed using
 43 the Schroeder method [7].

44 **De-biaser architecture** The de-biaser G takes a magnitude spectrogram as input. This is passed to a
 45 bi-directional LSTM with input size 257 and two hidden layers each of size 200, which produces an
 46 output with the same temporal length as the input spectrogram. This is passed through a linear layer
 47 of size 300 and a leakyReLU activation, followed by another linear layer of size 257 and a Sigmoid
 48 activation. The final mask is multiplied with the input magnitude spectrogram to create the generated
 49 magnitude spectrogram. A resynthesis module computes phase information from the input audio
 50 waveform, combines this with the generated magnitude spectrogram, and performs an inverse STFT
 51 to produce the generated waveform. The discriminator D consists of 4 2D Convolutional layers with
 52 kernel size (5,5) and 15 output channels, followed by a channel averaging operation and two linear
 53 layers of sizes 50 and 10. A LeakyReLU activation with negative slope = 0.3 is used after each
 54 intermediate layer. The final layer outputs a scalar-valued metric score estimate.

55 **De-biaser training** In stage (1) (see Sec. 3 — "Training" of our main paper), we train with batch
 56 size 32. During stage (3) fine-tuning, we use a batch size of 2. G and D are trained with learning
 57 rates of $2e-6$ and $5e-4$ respectively in both stages. In each epoch, We train on 10k samples randomly
 58 selected from the train set without replacement. The reverberator models R_v and R_b are updated with
 59 the target networks at a frequency of $E = 8$ epochs. For all models, we clip each audio sample to
 60 2.56s during training and evaluation.

61 **Reverberator training** We train the reverberators with batch size 4 and a learning rate of $1e-2$
 62 in stage (2). During stage (3) fine-tuning, we use batch size 2 and a learning rate of $1e-6$. Both
 63 reverberator models and the ViGAS baseline are trained with MSE loss between the log magnitude
 64 spectrogram of predicted and ground truth audio.

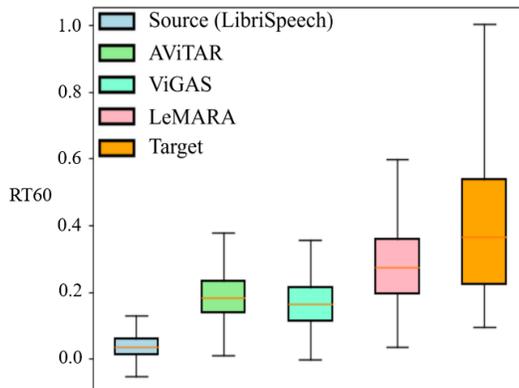


Figure 2: **LibriSpeech evaluation.** A more detailed version of Figure 4 (see main paper) with both baseline models for reference. The distribution of RT60 values for LeMARA reverberated audio (pink) better matches the ground truth distribution (orange) than either baseline (AViTAR and ViGAS).

65 **Baseline training details** We use a learning rate of $1e-2$ and a batch size of 4 to train ViGAS. We
 66 train AViTAR with batch size 4 — all other hyperparameters are set as described in [1].

67 **Compute** All models are trained on 8 NVIDIA Quadro RTX 6000 GPUs.

68 1.5 Augmentation strategy

69 We follow a data augmentation strategy similar to that proposed in [1] for training the baseline
 70 models, which was shown to produce better generalization performance on the LibriSpeech setting
 71 than when trained without this augmentation strategy. In particular, to each batch of dereverberated
 72 audio we add colored noise, perform a polarity inversion on the waveform with $p = 0.5$, and convolve
 73 the waveform with a randomly selected Room Impulse Response (RIR) from a different acoustic
 74 environment with $p = 0.9$. At test time, we evaluate without these audio augmentations. This strategy
 75 is designed to mask over residual acoustic information in dereverberated audio during training. We
 76 do not use this augmentation strategy in our approach as our model directly learns to remove residual
 77 acoustic information, obviating the need for a heuristic strategy to mask it out.

78 1.6 Limitations

79 Our approach focuses on visual acoustic matching on mono-channel audio exclusively. However,
 80 binaural cues in audio play a fundamental role in our perception of reverberation and room acoustics
 81 [3]. We leave it to future work to extend our approach to binaural audio.

82 1.7 Broader impact

83 While training on in-the-wild web videos allows wider access to a diverse variety of speakers and
 84 environments, it also introduces uncontrolled biases, speaker privacy concerns, and potentially
 85 harmful content into the model.

86 1.8 Data examples

87 Refer to video to view samples from both SoundSpaces-Speech and AVSpeech-Rooms.

88 References

89 [1] Changan Chen, Ruohan Gao, Paul Calamia, and Kristen Grauman. Visual acoustic matching. In
 90 *CVPR*, 2022.

- 91 [2] Changan Chen, Unnat Jain, Carl Schissler, Sebastia Vicenc Amengual Gari, Ziad Al-Halah,
92 Vamsi Krishna Ithapu, Philip Robinson, and Kristen Grauman. SoundSpaces: Audio-visual
93 navigation in 3d environments. In *ECCV*, 2020.
- 94 [3] Sasha Devore and Bertrand Delgutte. Effects of Reverberation on the Directional Sensitivity of
95 Auditory Neurons across the Tonotopic Axis: Influences of Interaural Time and Level Differences.
96 *Journal of Neuroscience*, 30(23):7826–7837, 2010. Publisher: Society for Neuroscience _eprint:
97 <https://www.jneurosci.org/content/30/23/7826.full.pdf>.
- 98 [4] Tiago H. Falk, Chenxi Zheng, and Wai-Yip Chan. A non-intrusive quality and intelligibility
99 measure of reverberant and dereverberated speech. *IEEE Transactions on Audio, Speech, and*
100 *Language Processing*, 18(7):1766–1774, 2010.
- 101 [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
102 recognition. In *CVPR*, pages 770–778, 2016.
- 103 [6] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without
104 convolution or region supervision, 2021.
- 105 [7] Manfred R. Schroeder. New method of measuring reverberation time. In *The Journal of the*
106 *Acoustical Society of America* 37, 409, 1965.