

SEMI-SUPERVISED SOURCE LOCALIZATION WITH DEEP GENERATIVE MODELING

Michael J. Bianco¹, Sharon Gannot², and Peter Gerstoft¹

¹ Marine Physical Laboratory, UCSD, CA, USA, mbianco@ucsd.edu

² Faculty of Engineering, Bar-Ilan University, Israel

ABSTRACT

We propose a semi-supervised localization approach based on deep generative modeling with variational autoencoders (VAE). Localization in reverberant environments remains a challenge, which machine learning (ML) has shown promise in addressing. Even with large data volumes, the number of labels available for supervised learning in reverberant environments is usually small. We address this issue by perform semi-supervised learning (SSL) with convolutional VAEs. The VAE is trained to generate the phase of relative transfer functions (RTFs), in parallel with a DOA classifier, on both labeled and unlabeled RTF samples. The VAE-SSL approach is compared with SRP-PHAT and fully-supervised CNNs. We find that VAE-SLL can outperform both SRP-PHAT and CNN in label-limited scenarios.

Index Terms— Source localization, semi-supervised learning, generative modeling, deep learning

1. INTRODUCTION

Source localization is an important problem in acoustics and many related fields. The performance of localization algorithms is degraded by reverberation, which induces complex temporal arrival structure at sensor arrays. Despite recent advances, e.g. [1–3], acoustic localization in reverberant environments remains a major challenge. There has been great interest in machine learning (ML)-based techniques in acoustics, including source localization and event detection [4–11]. A challenge for ML-based methods in acoustics is the limited amount of labeled data and the complex acoustic propagation in natural environments, despite large volumes of recordings [1, 2]. This limitation has motivated recent approaches for localization based on semi-supervised learning (SSL) [12, 13]. In SSL, ML models are trained using many examples with only few labels, with goal of exploiting the natural structure of the data.

We propose an SSL localization approach based on deep generative modeling with variational autoencoders

(VAEs) [14]. Deep generative models [15], e.g. generative adversarial networks (GANs) [16], have received much attention for their ability to learn high-dimensional sample distributions, including those of natural images [17]. GANs in acoustics have had success in generating raw audio [18] and speech enhancement [19]. VAEs learn from unlabeled data explicit latent codes for generating samples, and are inspiring examples of representation learning [20, 21].

We perform SSL based on VAEs to encode the phase of the relative transfer function (RTF) [22] between two microphones to a latent parametric distribution. The resulting model estimates DOA and generates RTF phase. By learning to generate RTF phase, the VAE-SSL system learns implicit physics from unlabeled samples. The VAE-SSL method is implemented using convolutional neural networks (CNNs). The performance of the convolutional VAE-SSL in reverberant environments is assessed against the steered response power with phase transform (SRP-PHAT) [23], and supervised convolutional neural network (CNN) approaches.

2. THEORY

We use RTFs [22], specifically the RTF phase, as the acoustic feature for our VAE-SSL approach. The RTF is independent of the source and well-represents the physics of the acoustic system. This helps to focus ML on physically relevant features. We encode the RTF phase as a function of source azimuth (direction of arrival, DOA).

2.1. Relative transfer function (RTF)

We consider time domain acoustic recordings of the form

$$d_i = a_i * s + u_i, \quad (1)$$

with s the source signal a_i ($i = \{1, 2\}$ the microphone index) the impulse responses (IRs) relating the source and each of the microphones, u_i noise signals which are independent of the source. Define the acoustic transfer functions $A_i(k)$ as the Fourier transform of the IRs a_i . Then, the relative transfer function (RTF) is defined as [22]

$$H(k) = \frac{A_2(k)}{A_1(k)}, \quad (2)$$

This work is supported by the Office of Naval Research, Grant No. N00014-11-1-0439 and the European Union’s Horizon 2020 Research and Innovation Programme under Grant Agreement No. 871245.

with k the frequency index. With d_1 as reference, $H(k)$ is

$$\hat{H}(k) = \frac{S_{d_2 d_1}}{S_{d_1 d_1}}, \quad (3)$$

with $S_{d_1 d_1} = D_1(k)^* D_1(k)$ the power spectral density (PSD), $S_{d_2 d_1} = D_2(k)^* D_1(k)$ the cross-PSD for a single frame and $*$ the complex conjugate. This estimator is biased since we neglect the PSD of the noise. An unbiased estimator can be obtained, but we observe here that the biased estimate $\hat{H}(k)$ works well, see also [24, 25]. For each FFT frame, a vector of RTFs is obtained $\hat{\mathbf{h}} = [\hat{H}(1) \dots \hat{H}(K)]^T \in \mathbb{C}^K$ for K frequency bins. We use RTFs estimated using a single frame as input features to the VAE-SSL.

The n th input sample to VAE-SSL and the supervised CNN is a sequence of RTF frames $\mathbf{x}_n = \text{vec}(\text{phase}(\hat{\mathbf{H}}_n)) \in \mathbb{R}^{KP}$, with $\hat{\mathbf{H}}_n = [\hat{\mathbf{h}}_n \dots \hat{\mathbf{h}}_{n+P-1}] \in \mathbb{C}^{K \times P}$, $K = N_{\text{FFT}}/2$, and P the number of RTF frames in the sequence.

2.2. Semi-supervised learning with VAEs

We assume the wrapped RTF phase \mathbf{x} is generated (disregarding subscript n to simplify notation) by a random process involving the latent random variable $\mathbf{z} \in \mathbb{R}^M$ and source location label $\mathbf{y} \in \mathbb{R}^T$ with T the number of candidate DOAs. We formulate a principled semi-supervised learning framework based on VAEs, which treats the label \mathbf{y} as both latent and observed, and trains a classifier using both labeled and unlabeled data. This corresponds to the ‘M2’ model in [20]. Starting with Bayes’ rule we have for labeled data

$$p_\theta(\mathbf{z}|\mathbf{x}, \mathbf{y}) = \frac{p_\theta(\mathbf{x}, \mathbf{y}|\mathbf{z})p_\theta(\mathbf{z})}{p_\theta(\mathbf{x}, \mathbf{y})} \quad (4)$$

and for unlabeled data

$$p_\theta(\mathbf{y}, \mathbf{z}|\mathbf{x}) = \frac{p_\theta(\mathbf{x}|\mathbf{y}, \mathbf{z})p_\theta(\mathbf{y}, \mathbf{z})}{p_\theta(\mathbf{x})} \quad (5)$$

Using (4) as an example, direct estimation of the posterior $p_\theta(\mathbf{z}|\mathbf{x}, \mathbf{y})$ is nearly always intractable due to $p(\mathbf{x}, \mathbf{y}) = \int p(\mathbf{x}, \mathbf{y}, \mathbf{z})d\mathbf{z}$. As will later be shown, the subscript θ indicates distributions and functions defined using the decoder network, i.e. the generative model in the VAE.

VAEs [20, 26] approximate posterior distributions using variational inference (VI) [27], a family of methods for approximating conditional densities which relies on optimization instead of (MCMC) sampling. In VAEs the conditional densities are modeled with NNs. A variational approximation to the intractable posterior $p_\theta(\mathbf{z}|\mathbf{x}, \mathbf{y})$ is defined by the encoder network as $q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y}) \approx p_\theta(\mathbf{z}|\mathbf{x}, \mathbf{y})$, with the subscript ϕ corresponding to distributions and functions defined using the encoder network. The networks constituting the VAE-SSL model are shown in Fig. 1.




Fig. 1: Directed graphical model for semi-supervised learning with VAE. (a) Inference models (encoders), supervised learning (labeled) per (8),(12). (b) Inference models, unsupervised (unlabeled) per (10). (c) Generative model (decoder).

Starting with the model for the labeled data (see (4)), per VI we seek $q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y})$ which minimizes the KL-divergence

$$\{\phi, \theta\} = \arg \min_{\phi, \theta} \text{KL}(q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y}) || p_\theta(\mathbf{z}|\mathbf{x}, \mathbf{y})). \quad (6)$$

Considering first the labeled data, the intractable posterior is approximated by $q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y}) \approx p_\theta(\mathbf{z}|\mathbf{x}, \mathbf{y})$. Assessing the KL-divergence, we obtain

$$\begin{aligned} \text{KL}(q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y}) || p_\theta(\mathbf{z}|\mathbf{x}, \mathbf{y})) &= \mathbb{E}[\log q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y})] - \mathbb{E}[\log p_\theta(\mathbf{x}, \mathbf{y}|\mathbf{z})p_\theta(\mathbf{z})] \\ &\quad + \log p_\theta(\mathbf{x}, \mathbf{y}) \\ &= -\text{ELBO} + \log p_\theta(\mathbf{x}, \mathbf{y}), \end{aligned} \quad (7)$$

with the expectation \mathbb{E} relative to $q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y})$. This reveals the dependence of the KL divergence on evidence $p_\theta(\mathbf{x}, \mathbf{y})$, which is intractable. The other two terms in (7) form the evidence lower bound (ELBO). Since the KL divergence is non-negative, the ELBO ‘lower bounds’ the evidence: $\text{ELBO} \leq \log p_\theta(\mathbf{x}, \mathbf{y})$. Maximizing the ELBO is equivalent to minimizing the KL (6). We thus minimize $-\text{ELBO}$.

Considering the ELBO terms from (7), we formulate the objective for labeled data, assuming \mathbf{y} and \mathbf{z} independent

$$\begin{aligned} -C(\theta, \phi; \mathbf{x}, \mathbf{y}) &= \mathbb{E}[\log p_\theta(\mathbf{x}, \mathbf{y}|\mathbf{z})p_\theta(\mathbf{z}) - \log q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y})] \\ &= \mathbb{E}[\log p_\theta(\mathbf{x}|\mathbf{y}, \mathbf{z}) + \log p_\theta(\mathbf{y}) + \log p_\theta(\mathbf{z}) \\ &\quad - \log q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{y})]. \end{aligned} \quad (8)$$

Next, an objective for unlabeled data is derived. The intractable posterior $q_\phi(\mathbf{y}, \mathbf{z}|\mathbf{x}) \approx p_\theta(\mathbf{y}, \mathbf{z}|\mathbf{x})$ from (5). From the KL, we find the objective (negative ELBO) as

$$\begin{aligned} -D(\theta, \phi; \mathbf{x}) &= \mathbb{E}[\log p_\theta(\mathbf{x}|\mathbf{y}, \mathbf{z})p_\theta(\mathbf{y}, \mathbf{z}) - \log q_\phi(\mathbf{z}, \mathbf{y}|\mathbf{x})] \\ &= \mathbb{E}[\log p_\theta(\mathbf{x}|\mathbf{y}, \mathbf{z}) + \log p_\theta(\mathbf{y}) + \log p_\theta(\mathbf{z}) \\ &\quad - \log q_\phi(\mathbf{y}|\mathbf{x}) - \log q_\phi(\mathbf{z}|\mathbf{x})] \end{aligned} \quad (9)$$

Table 1: Design case: MAE (°) and Accuracy (%) of VAE-SSL and alternative approaches. SRP-PHAT results in bottom row.

J (# labels)	VAE-SSL		CNN	
	MAE	Acc.	MAE	Acc.
37	27.2	9.59	39.5	13.4
74	10.7	36.0	33.3	21.3
148	13.7	39.5	27.3	24.7
481	4.37	69.8	8.32	59.4
999	2.24	83.7	0.34	96.4
SRP-PHAT	18.0	11.6		

Table 2: Validation case: Notation as Table 1.

J (# labels)	VAE-SSL		CNN	
	MAE	Acc.	MAE	Acc.
37	29.2	9.03	49.6	10.2
74	14.5	31.5	43.8	17.0
148	17.2	36.2	34.6	21.4
481	7.00	65.9	10.2	56.9
999	5.48	74.0	1.32	90.5
SRP-PHAT	21.7	12.3		

with the expectation relative to $q_\phi(\mathbf{z}, \mathbf{y}|\mathbf{x})$. Further expanding (9) we obtain

$$\begin{aligned}
 -D(\theta, \phi; \mathbf{x}) &= \mathbb{E}_{q_\phi(\mathbf{y}|\mathbf{x})} [\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{y}, \mathbf{z}) + \log p_\theta(\mathbf{y}) \\
 &\quad + \log p_\theta(\mathbf{z}) - \log q_\phi(\mathbf{y}|\mathbf{x}) - \log q_\phi(\mathbf{z}|\mathbf{x})]] \\
 &= \sum_{\mathbf{y}} q_\phi(\mathbf{y}|\mathbf{x}) [-C(\theta, \phi; \mathbf{x}, \mathbf{y}) - \log q_\phi(\mathbf{y}|\mathbf{x})] \quad (10)
 \end{aligned}$$

An overall objective for training using labelled and unlabelled data is derived by combining (8),(10) as

$$\mathcal{J} = \sum_{\{\mathbf{x}, \mathbf{y}\}} C(\theta, \phi; \mathbf{x}, \mathbf{y}) + \sum_{\{\mathbf{x}\}} D(\theta, \phi; \mathbf{x}). \quad (11)$$

Assessing the terms in (8), the supervised learning objective C does not condition the label \mathbf{y} on the sample \mathbf{x} . The term $\log q_\phi(\mathbf{y}|\mathbf{x})$ is only present in the unsupervised learning objective (11). This issue is remedied by an additional term

$$\mathcal{J}^\alpha = \mathcal{J} + \alpha \sum_{\{\mathbf{x}, \mathbf{y}\}} -\log q_\phi(\mathbf{y}|\mathbf{x}). \quad (12)$$

It is assumed that the data is explained by the generative process (see (10) for terms): $p_\theta(\mathbf{y}) = \text{Cat}(\mathbf{y}|\boldsymbol{\pi})$, with $\text{Cat}(\cdot|\cdot)$ the categorical (multinomial) distribution and $\boldsymbol{\pi} \in \mathbb{R}^T$ the probabilities of the classes (which are assumed equal, and normalized such that $\sum_T \pi_t = 1$); $p_\theta(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$ (as before); and $p_\theta(\mathbf{x}|\mathbf{y}, \mathbf{z}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_\theta(\mathbf{y}, \mathbf{z}), \mathbf{I})$, with $\boldsymbol{\mu}_\theta(\mathbf{y}, \mathbf{z})$ the decoder. The densities of the inference model are: $q_\phi(\mathbf{z}|\mathbf{y}, \mathbf{x}) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_\phi(\mathbf{x}, \mathbf{y}), \text{diag}(\boldsymbol{\sigma}_\phi^2(\mathbf{x}, \mathbf{y})))$, with $\boldsymbol{\mu}_\phi(\mathbf{x}, \mathbf{y})$ and $\boldsymbol{\sigma}_\phi^2(\mathbf{x}, \mathbf{y})$ the outputs of the z-encoder; and $q_\phi(\mathbf{y}|\mathbf{x}) = \text{Cat}(\mathbf{y}|\boldsymbol{\pi}_\phi(\mathbf{x}))$, with $\boldsymbol{\pi}_\phi(\mathbf{x})$ the classifier network.

Table 3: Test case I: Notation as Table 1.

J (# labels)	VAE-SSL		CNN	
	MAE	Acc.	MAE	Acc.
37	27.3	10.3	43.9	10.6
74	11.0	33.2	37.2	17.6
148	13.7	37.1	29.2	22.4
481	5.78	66.0	10.5	54.1
999	2.97	76.8	0.57	93.2
SRP-PHAT	20.2	10.1		

Table 4: Test case II: Notation as Table 1.

J (# labels)	VAE-SSL		CNN	
	MAE	Acc.	MAE	Acc.
37	28.5	9.63	50.0	8.90
74	13.0	31.7	43.5	15.3
148	16.0	36.6	33.2	21.0
481	6.19	65.9	11.5	53.8
999	4.28	74.4	0.85	92.1
SRP-PHAT	22.0	10.3		

Thus in this case, we have 3 networks: 1) the label inference (classifier) network $\boldsymbol{\pi}_\phi(\mathbf{x})$ corresponding to $q_\phi(\mathbf{y}|\mathbf{x})$; 2) the inference network $\boldsymbol{\mu}_\phi(\mathbf{x}, \mathbf{y})$ and $\boldsymbol{\sigma}_\phi^2(\mathbf{x}, \mathbf{y})$ corresponding to $q_\phi(\mathbf{z}|\mathbf{y}, \mathbf{x})$; and 3) the decoder (generative) network $\boldsymbol{\mu}_\theta(\mathbf{y}, \mathbf{z})$ corresponding to $p_\theta(\mathbf{x}|\mathbf{y}, \mathbf{z})$. Graphical models representing the inference and generative networks are shown in Fig. 1. For optimization of the VAE-SSL system, the objective (12) is evaluated using probabilistic programming [28].

3. EXPERIMENTS

We assess the DOA estimation performance of the VAE-SSL approaches in moderately reverberant environments against two alternative techniques: SRP-PHAT [29], and a supervised CNN baseline. We simulate 4 room configurations to test the generalization of the learning-based methods in label-limited scenarios, as we expect this to closely approximate real applications. Thus, we test the learning-based methods (VAE-SSL and CNN) using few supervised labels, with J the number of labels. The results are summarized in Tables 1–4, giving the performance of each method in terms of DOA error (mean absolute error, MAE) and frame-level accuracy.

All algorithms were coded in Python and the reverberant data was generated using Matlab [30]. The VAE-SSL system and CNN were implemented using Pytorch [31], with the Pyro package [28] used for stochastic variational inference and optimization of the VAE-SSL. The NNs were optimized using Adam.




Fig. 2: Normalized histograms of DOA estimates by VAE-SSL and SRP-PHAT for each DOA index, for *Test case II*. VAE-SSL (a) $J = 481$ and (b) $J = 148$ labels. (c) SRP-PHAT results. True DOA is shown in each plot as white dashed line.

3.1. Reverberant room data

The reverberant acoustic data were generated using the Room Impulse Response (RIR) generator [30]. The synthetic IRs are convolved with a Gaussian time-domain signal $s \sim \mathcal{N}(0, 1)$ (see (1)) and FFTs and RTFs were obtained for two sensors with only one active source location for each time bin. Noise signals were used for training/testing to avoid the difficulty of voice activity detection. We simulated 4 room configurations to test the generalization of the learning-based methods in label-limited scenarios. This included off-design conditions with changes to reverberation time and perturbations in the microphone positions.

The nominal room configuration, deemed the ‘design’ case (see Table 1 for results) is a square room with x-y-z dimensions $6 \times 6 \times 2.4$ m with a reverberation time $RT_{60} = 500$ ms and $c = 343$ m/s. Two omnidirectional microphones were located in the center of the room with spacing 8 cm. The source range was 1.5 m, with sources at 5° resolution for $[-90^\circ, 90^\circ]$ azimuth relative to the array broadside ($T = 37$ candidate DOAs). We used 1 s signals for each sensor location to obtain FFT and RTF frames with a 16 kHz sampling rate. A sensor noise level of 20dB was assumed and 10 signal realizations were generated for each candidate DOA.

Three off-design rooms were simulated. The first, deemed

the *validation* case (see Table 2) had the same room configuration as the design case, but the reverberation time $RT_{60} = 700$ ms. The second, deemed *test case I* (see Table 3), had the room configuration and reverberation time as the design case, but the two microphone locations were displaced by 5 mm and 3 mm in opposite directions along the y-axis. The third, deemed *test case II* (see Table 4), had the same room configuration as test case I, but $RT_{60} = 600$ ms.

The signal at the microphones are given in (1). We obtain the RTFs from the data by (3). The RTFs are estimated using single FFT frames with Hamming windowing with 50% overlap and $N_{FFT} = 256$. The VAE and supervised CNN inputs \mathbf{x}_n use $P = 32$ RTF vectors, giving an input size $K \times P = 32 \times 128$. For fair comparison, SRP-PHAT used the $P = 32$ FFT snapshots used to estimate the RTFs. Thus, the length of the analysis frame for all the methods was 0.26 s. This yielded $\sim 44,000$ (RTF) frames per room configuration.

3.2. Learning-based model parameters

The VAE-SSL model (classifier, inference, and generative networks) were implemented using CNNs. The classifier network $\pi_\phi(\mathbf{x})$ and inference network ($\mu_\phi(\mathbf{x}, \mathbf{y})$ and $\sigma_\phi^2(\mathbf{x}, \mathbf{y})$) used the same architecture: two convolutional layers with 3-by-3 kernels with 2-by-2 max pooling, followed by two fully-connected layers with 200 hidden units and output units. Each of the convolutional layers had 8 channels. The hidden units and convolutional layers used ReLU activation. The classifier network had T output units with softmax activation. The inference network had $2 \times M$ output units, corresponding to the latent code with dimension M , with linear activation. The convolution operations were applied to the RTF frames \mathbf{x} , and the labels \mathbf{y} were input to the graph at the 200-unit hidden layer. The latent code dimension for all experiments was $M = 2$, assuming that the hidden representation must account for source range (near-field) and temporal variation of the RTFs.

The generative network ($\mu_\theta(\mathbf{y}, \mathbf{z})$) mirrors the inference network architecture, and consists of two fully connected layers of 200 and 2048 hidden units, followed by two transpose convolutional layers, each with a 2-by-2 max-unpool operations and transpose convolutions with 3-by-3 kernels. The hidden units and transpose convolution layer used ReLU activation. The output of the final layer had linear activation. Like the inference network, the labels \mathbf{y} were input to the graph at the 200-unit hidden layer.

For comparison, we used the baseline ‘supervised-only’ CNN from the VAE-SSL model, without unsupervised learning and integration with the generative model.

For all cases, the learning rate was $5e-5$ and the batch size was 256. The default momentum and decay values (0.9, 0.999) were used from the Pytorch implementation of Adam. The VAE-SSL has the tuning parameter, the auxiliary multiplier α , which weights the supervised objective (see

(12)). It was found in the experiments that the performance of VAE-SSL was not very sensitive to α , though the best value per validation accuracy was found by grid search.

3.3. Non-learning: SRP-PHAT configuration

The FFT features (used to calculate RTF) were used for the SRP-PHAT approach with the same number of frames as the VAE-SSL and CNN methods. SRP-PHAT used 37 candidate DOAs over $[-90^\circ, 90^\circ]$. SRP-PHAT was implemented using the Pyroomacoustics toolbox [32].

3.4. Training and performance

For VAE-SSL and supervised CNN training, J labelled examples were drawn from the *design* case data to train the networks (i.e. calculate the losses and gradients) and J labelled examples were drawn from the *validation* case data for validation. The model was chosen based on validation accuracy. For VAE-SSL, the additional unlabeled examples were used to train the networks.

The full set of RTF features and labels for the experiments are $\{\mathbf{x}_n, \mathbf{y}_n\} \in \mathcal{D}$. From this set supervised samples with $\{\mathbf{x}_j, \mathbf{y}_j\} \in \mathcal{S}$ with j the supervised sample index. The unsupervised examples (unlabeled) are $\{\mathbf{x}_i\} \notin \mathcal{S}$. All the RTF frames $\mathbf{X} = [\mathbf{x}_1 \dots \mathbf{x}_N]$ were normalized to the range $[0,1]$.

We consider a range of J for assessing the performance of the DOA algorithms. The performance on the design, validation, test I, and test II cases (see Table 1–4) was assessed on the sample labels not available in training. We use $J = 37, 74, 148, 481, 999$, which are multiples of the number of candidate DOAs ($T = 37$) to ensure an equal number of labelled samples for each DOA. The case where $J = 37$, corresponds to one labelled input frame for each DOA. For $J = 999$, there are 27 labelled input frames available per DOA.

In VAE-SSL, for each value of J , the remaining samples are used for unsupervised learning, assuming no labels are available for those samples. During each training epoch, the supervised batches are used at a frequency proportional to their proportion of the overall data (supervised and unsupervised). Only the supervised samples were used to train the supervised CNN. Since for the unlabeled examples, the DOAs are technically unavailable, there are unsupervised frames \mathbf{x}_i which contain a DOA change. For the supervised examples only input frames with a single DOA are used.

The VAE-SSL auxiliary multiplier α values for each value of J were chosen by grid search in the interval $[10,100]$ in increments of 10. For $J = 37, 74, 148, 481, 999$, the corresponding $\alpha = 10, 50, 70, 20, 80$.

The performance of the VAE-SSL method generally exceeds the supervised CNN and SRP-PHAT by a large margin for these label-limited scenarios, both in terms of MAE and accuracy. It is robust to small perturbations in sensor location and reverberation levels different from training. The

supervised CNN outperforms the VAE-SSL approach for the maximum number of labels ($J = 999$). However, as the number of labels decreases, the performance of VAE-SSL exceeds that of the supervised CNN and SRP-PHAT. VAE-SSL outperforms SRP-PHAT, both in terms of MAE and accuracy for as few as two labeled input samples for each DOA ($J = 74$). The performance of VAE-SSL is compared with SRP-PHAT in Fig. 2 for only few labelled examples. In general, SRP-PHAT was outperformed by the learning-based methods.

4. CONCLUSIONS

This study shows that deep generative modeling can localize sources well in reverberant environments, relative to existing approaches, when only few labels are available. It was demonstrated for moderately reverberant scenarios that VAE-SSL outperforms supervised CNN when only few labels are available, provided significant unlabeled data is available. Training the VAE-SSL on all reverberant frames (with and without labels) allows the VAE system to exploit the structure of the data. The experiments show, only two labeled samples per DOA permit the VAE-SSL to obtain better performance than SRP-PHAT.

5. REFERENCES

- [1] H. Purwins, B. Li, T. Virtanen, J. Schluter, S.-Y. Chang, and T. Sainath, "Deep learning for audio signal processing," *IEEE J. Sel. Topics Signal Process.*, vol. 13, no. 2, pp. 206–219, May 2019.
- [2] M. J. Bianco, P. Gerstoft, J. Traer, E. Ozanich, M. A. Roch, S. Gannot, and C.-A. Deledalle, "Machine learning in acoustics: Theory and applications," *J. Acoust. Soc. Am.*, vol. 146, no. 5, pp. 3590–3628, 2019.
- [3] S. Gannot, M. Haardt, W. Kellermann, and P. Willett, "Introduction to the issue on acoustic source localization and tracking in dynamic real-life scenes," *IEEE J. Sel. Topics Signal Process.*, vol. 13, no. 1, pp. 3–7, 2019.
- [4] H. Nakashima and T. Mukai, "3D sound source localization system based on learning of binaural hearing," in *IEEE Int. Conf. Syst., Man, Cybern. IEEE*, 2005, vol. 4, pp. 3534–3539.
- [5] A. Deleforge and R. Horaud, "2D sound-source localization on the binaural manifold," in *IEEE Int. Workshop Mach. Learn. Signal Process.* IEEE, 2012, pp. 1–6.
- [6] A. Mesaros, T. Heittola, A. Diment, B. Elizalde, A. Shah, E. Vincent, B. Raj, and T. Virtanen, "Dcase 2017 challenge setup: Tasks, datasets and baseline system," in *DCASE 2017*, 2017.
- [7] E. Vincent, T. Virtanen, and S. Gannot, *Audio source separation and speech enhancement*, John Wiley & Sons, 2018.
- [8] S. Chakrabarty and E. A. P. Habets, "Multi-speaker doa estimation using deep convolutional networks trained with noise signals," *IEEE J. Sel. Topics Signal Process.*, vol. 13, no. 1, pp. 8–21, Mar. 2019.

- [9] S. Adavanne, A. Politis, J. Nikunen, and T. Virtanen, "Sound event localization and detection of overlapping sources using convolutional recurrent neural networks," *IEEE J. Sel. Topics Signal Process.*, 2019.
- [10] G. Ping, E. Fernandez-Grande, P. Gerstoft, and Z. Chu, "Three-dimensional source localization using sparse bayesian learning on a spherical microphone array," *J. Acoust. Soc. Am.*, vol. 147, no. 6, pp. 3895–3904, 2020.
- [11] E. Ozanich, P. Gerstoft, and H. Niu, "A feedforward neural network for direction-of-arrival estimation," *J. Acoust. Soc. Am.*, vol. 147, no. 3, pp. 2035–2048, 2020.
- [12] B. Laufer-Goldshtein, R. Talmon, and S. Gannot, "Semi-supervised sound source localization based on manifold regularization," *IEEE/ACM Trans. Audio, Speech, and Lang. Process.*, vol. 24, no. 8, pp. 1393–1407, Aug. 2016.
- [13] R. Opochninsky, B. Laufer-Goldshtein, S. Gannot, and G. Chechik, "Deep ranking-based sound source localization," in *IEEE Workshop Appl. Signal Process. Audio Acoustic. (WASPAA)*, 2019, pp. 283–287.
- [14] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *Proc. Int. Conf. Learn. Represent.*, 2014.
- [15] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep learning*, vol. 1, MIT press Cambridge, 2016.
- [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Adv. Neural Info. Process. Sys.*, 2014, pp. 2672–2680.
- [17] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of stylegan," *arXiv preprint arXiv:1912.04958*, 2019.
- [18] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016.
- [19] C. Donahue, B. Li, and R. Prabhavalkar, "Exploring speech enhancement with generative adversarial networks for robust speech recognition," in *IEEE Int. Conf. Acoust. Speech and Signal Process. (ICASSP)*, 2018, pp. 5024–5028.
- [20] D. P. Kingma, S. Mohamed, D. J. Rezende, and M. Welling, "Semi-supervised learning with deep generative models," in *Adv. Neural Info. Process. Sys.*, 2014, pp. 3581–3589.
- [21] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: a review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, Aug. 2013.
- [22] S. Gannot, D. Burshtein, and E. Weinstein, "Signal enhancement using beamforming and nonstationarity with applications to speech," *IEEE Trans. Signal Process.*, vol. 49, no. 8, pp. 1614–1626, Aug. 2001.
- [23] M. S. Brandstein and H. F. Silverman, "A robust method for speech signal time-delay estimation in reverberant rooms," in *IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, 1997, vol. 1, pp. 375–378.
- [24] S. Markovich-Golan, S. Gannot, and W. Kellermann, "Performance analysis of the covariance-whitening and the covariance-subtraction methods for estimating the relative transfer function," in *EUSIPCO*, 2018, pp. 2499–2503.
- [25] Z. Koldovský, J. Málek, and S. Gannot, "Spatial source subtraction based on incomplete measurements of relative transfer function," *IEEE/ACM Trans. Audio, Speech, and Lang. Process.*, vol. 23, no. 8, pp. 1335–1347, 2015.
- [26] D. P. Kingma, M. Welling, et al., "An introduction to variational autoencoders," *Found. and Trends in Machine Learn.*, vol. 12, no. 4, pp. 307–392, 2019.
- [27] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, "Variational inference: A review for statisticians," *J. Amer. Stat. Assoc.*, vol. 112, no. 518, pp. 859–877, 2017.
- [28] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, . Karaletsos, R. Singh, P. Szerlip, P. Horsfall, and N. D. Goodman, "Pyro: Deep Universal Probabilistic Programming," *J. Mach. Learn. Res.*, 2018.
- [29] J.H DiBiase, H.F Silverman, and M.S Brandstein, "Robust localization in reverberant rooms," in *Microphone Arrays*, pp. 157–180. Springer, 2001.
- [30] E. A. P. Habets, "Room Impulse Response (RIR) generator," (Online) Available: <https://github.com/ehabets/RIR-Generator>, 2016.
- [31] A. Paszke, S. Gross, F. Massa, A. Lerer, and J. Bradbury et al., "Pytorch: An imperative style, high-performance deep learning library," in *Adv. Neural Info. Process. Sys.*, 2019, pp. 8024–8035.
- [32] R. Scheibler, E. Bezzam, and I. Dokmanić, "Pyroomacoustics: A python package for audio room simulation and array processing algorithms," in *IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, 2018, pp. 351–355.