

Latent Iterative Refinement for Modular Source Separation

Bralios, Dimitrios; Tzinis, Efthymios; Wichern, Gordon; Smaragdis, Paris; Le Roux, Jonathan

TR2023-019 April 14, 2023

Abstract

Traditional source separation approaches train deep neural network models end-to-end with all the data available at once by minimizing the empirical risk on the whole training set. On the inference side, after training the model, the user fetches a static computation graph and runs the full model on some specified observed mixture signal to get the estimated source signals. Additionally, many of those models consist of several basic processing blocks which are applied sequentially. We argue that we can significantly increase resource efficiency during both training and inference stages by reformulating a model's training and inference procedures as iterative mappings of latent signal representations. First, we can apply the same processing block more than once on its output to refine the input signal and consequently improve parameter efficiency. During training, we can follow a block-wise procedure which enables a reduction on memory requirements. Thus, one can train a very complicated network structure using significantly less computation compared to end-to-end training. During inference, we can dynamically adjust how many processing blocks and iterations of a specific block an input signal needs using a gating module.

*IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)
2023*

LATENT ITERATIVE REFINEMENT FOR MODULAR SOURCE SEPARATION

Dimitrios Bralios¹, Efthymios Tzinis¹, Gordon Wichern², Paris Smaragdis¹, Jonathan Le Roux²

¹University of Illinois Urbana-Champaign, Urbana, IL, USA

²Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

ABSTRACT

Traditional source separation approaches train deep neural network models end-to-end with all the data available at once by minimizing the empirical risk on the whole training set. On the inference side, after training the model, the user fetches a static computation graph and runs the full model on some specified observed mixture signal to get the estimated source signals. Additionally, many of those models consist of several basic processing blocks which are applied sequentially. We argue that we can significantly increase resource efficiency during both training and inference stages by reformulating a model’s training and inference procedures as iterative mappings of latent signal representations. First, we can apply the same processing block more than once on its output to refine the input signal and consequently improve parameter efficiency. During training, we can follow a block-wise procedure which enables a reduction on memory requirements. Thus, one can train a very complicated network structure using significantly less computation compared to end-to-end training. During inference, we can dynamically adjust how many processing blocks and iterations of a specific block an input signal needs using a gating module.

Index Terms— Iterative latent refinement, progressive training, sound source separation, layer-wise optimization

1. INTRODUCTION

Layer-wise optimization is a widely used idea in terms of progressively training a deep neural network using shallow layer-wise losses [1, 2]. This technique has been successfully applied towards applications such as image classification [3], image generation [2], or speech enhancement [4–7], and has been shown to perform comparably with end-to-end training but with large computational savings. This is especially important when one can only fit sub-parts of the forward and backward procedure in memory, e.g. a single trainable layer and not the full model.

Recursive neural networks are centered around the idea of weight-sharing in a sequence of layers [8–11]. Therefore, they have been proposed as a way to reduce the overall number of parameters of models. In the case of speech enhancement, [12] proposes a network consisting of multiple stages where feedback from the previous stage is provided. The concept of iterative refinement is also found in more recent diffusion based models [13, 14].

Further, dynamic neural networks and adaptive computation are other broad concepts that have been explored recently [15]. On the inference side, previous works have proposed to train models which are able to perform early-exit and save computation for inputs which are less noisy or sometimes when more denoising could lead to more

artifacts that would harm the performance of downstream tasks [4, 6, 16, 17]. An early exiting mechanism centered around a manually set threshold on a distance between the output of successive processing stages, is proposed for the problem of source separation [17] as well as for speech enhancement [18]. In other tasks, different approaches utilize a module that controls the flow of computation allowing data-adaptive computation. For instance, in [19], a gating module has been proposed in order to controls layer-wise skip connections and is trained using reinforcement learning methods. In [20], instead of skipping layers, the gating module controls the number of iterations for some layers, while in [21] the gating function is trained jointly.

In this work, we build upon these previous results and follow a more holistic approach towards utilizing both block-wise training (such as Bloom-Net [7]) and inference recipes, with the important addition that we can also perform multiple applications of the same blocks (with the same parameters), consequently reducing the model size. In the case where we progressively train the network, we enable a much larger parallelization by greatly reducing computational demands during training.

During the process of refinement, some of the noisy signals approach their ideal targets in a small number of steps even when applying the same layer multiple times. This indicates that a share of the samples can be processed in fewer steps without causing a significant drop in performance. To that end, we can transform the separation model into an adaptive network which determines the amount of computation needed on a sample-by-sample basis. We achieve this by attaching to the separation network a learnable gating module (based on [20]) that for each input decides when to take an early exit. Joint training of the separation model and the gating module not only renders the network dynamic but it also enables it to adapt to different inputs. Our experiments in both source separation and speech enhancement show the effectiveness of our proposed holistic latent iterative refinement approach.

2. LATENT ITERATIVE REFINEMENT

Given an input mixture \mathbf{x} , the objective of a source separation network is to recover the sources \mathbf{s} that compose it. A large class of source separation architectures fit the following definition:

$$\begin{aligned}\hat{\mathbf{m}} &= \text{Mask}[\mathcal{S}(\mathcal{E}(\mathbf{x}))], \\ \hat{\mathbf{s}} &= \mathcal{D}(\mathcal{E}(\mathbf{x}) \odot \hat{\mathbf{m}}),\end{aligned}\tag{1}$$

where \mathcal{E} is an encoder which transforms the time-domain input mixture into a latent representation. This latent representation is then fed into a separation module \mathcal{S} whose output is transformed using a Mask network into mask estimates $\hat{\mathbf{m}}$ which are applied on the mixture’s latent representation. Finally, \mathcal{D} is a decoder which turns the latent-space source estimates into the time-domain estimated sources $\hat{\mathbf{s}}$. Usually, the separation module \mathcal{S} is a composition of blocks which have the same structure, as shown in Fig. 1:

D. Bralios and P. Smaragdis were partially funded by NIFA grant #2020-67021-32799. E. Tzinis was partially funded by the Google Ph.D. fellowship.




Fig. 1: A separation module \mathcal{S} composed of blocks \mathcal{B}_{θ_i} , each of which refines a latent-space estimate.

$$\mathcal{S}(\mathbf{v}) = (\mathcal{B}_{\theta_n} \circ \mathcal{B}_{\theta_{n-1}} \circ \dots \circ \mathcal{B}_{\theta_1})(\mathbf{v}), \quad (2)$$

where \circ denotes composition, \mathcal{B}_{θ_i} denotes the i -th block with parameters θ_i , and $\mathbf{v} \in \mathbb{R}^{C \times L}$ is the latent mixture representation fed to the separation module. Formulating the separation module this way, we can regard the effect of the blocks as an iterative refinement process of a latent-space estimate. Thus, based on this formulation, we present a unified framework encompassing a series of techniques with multiple uses centered around resource efficiency.

2.1. Iterative Block-Wise Latent Refinement

Casting the separation module as a refinement process, we can impose the constraint of shared parameters between its blocks. In this case, we have $\theta_i := \theta, \forall i \in \{1, \dots, n\}$, and the separation module becomes as such:

$$\mathcal{S}(\mathbf{v}) = (\underbrace{\mathcal{B}_{\theta} \circ \mathcal{B}_{\theta} \circ \dots \circ \mathcal{B}_{\theta}}_{n \text{ times}})(\mathbf{v}), \quad (3)$$

where in essence the block \mathcal{B}_{θ} is applied iteratively to its output. This allows us to drastically reduce the number of model parameters albeit at the cost of a performance drop. In order to mitigate this reduction in performance, we can increase the size of a block \mathcal{B}_{θ_i} , which can be composed of k distinct sub-blocks $\mathcal{C}_{\theta_{i,j}}$ as follows:

$$\mathcal{B}_{\theta_i}(\mathbf{v}) = (\mathcal{C}_{\theta_{i,k}} \circ \mathcal{C}_{\theta_{i,k-1}} \circ \dots \circ \mathcal{C}_{\theta_{i,1}})(\mathbf{v}). \quad (4)$$

Further, we can have more than one iterated block. For example, given the blocks \mathcal{B}_{θ_1} and \mathcal{B}_{θ_2} , we can construct the following separation module:

$$\mathcal{S}(\mathbf{v}) = ((\mathcal{B}_{\theta_2} \circ \dots \circ \mathcal{B}_{\theta_2}) \circ (\mathcal{B}_{\theta_1} \circ \dots \circ \mathcal{B}_{\theta_1}))(\mathbf{v}), \quad (5)$$

which is depicted in Fig. 2. In the case of m distinct iterated blocks, the network is defined as:

$$\mathcal{S}(\mathbf{v}) = ((\mathcal{B}_{\theta_m} \circ \dots \circ \mathcal{B}_{\theta_m}) \circ \dots \circ (\mathcal{B}_{\theta_1} \circ \dots \circ \mathcal{B}_{\theta_1}))(\mathbf{v}). \quad (6)$$




Fig. 2: A separation module \mathcal{S} constructed from two different blocks \mathcal{B}_{θ_1} and \mathcal{B}_{θ_2} iterated n_1 and n_2 times respectively.

2.2. Progressive Latent Representation Training

Instead of learning the complete refinement process in an end-to-end manner, we can split up the training into m different steps, where m is the number of distinct blocks. As before, the separation module \mathcal{S}




Fig. 3: Separation module at step i of the progressive training where block \mathcal{B}_{θ_i} is trained while the previous iterated blocks $\mathcal{B}_{\theta_1}, \mathcal{B}_{\theta_2}, \dots, \mathcal{B}_{\theta_{i-1}}$ are kept frozen, as represented by the dotted lines. Feedback connections represent iterations depicted in a more compact way.

is given by Eq. 6 and each block \mathcal{B}_{θ_i} can be composed of several sub-blocks $\mathcal{C}_{\theta_{i,j}}$ as described in Eq. 4.

In the first training step, we learn the shared encoder \mathcal{E} and the first block \mathcal{B}_{θ_1} along with its dedicated mask network and decoder \mathcal{D}_1 . After this step is completed, we keep the parameters of the shared encoder \mathcal{E} and the first block \mathcal{B}_{θ_1} frozen as we proceed. At each of the following steps, a block \mathcal{B}_{θ_i} along with its dedicated decoder \mathcal{D}_i and its mask network are trained while the previous iterated blocks $\mathcal{B}_{\theta_1}, \mathcal{B}_{\theta_2}, \dots, \mathcal{B}_{\theta_{i-1}}$ (as shown in Fig. 3) and the shared encoder \mathcal{E} are kept fixed.

This technique allows us to reduce the computational requirements of the training procedure in terms of memory. Additionally, after training is completed, we have m different models of increasing size and computational requirements which can be deployed depending on the available resources.

2.3. Adaptive Early Exit

Instead of having a static refinement process, we can turn it into a dynamic process where the number of processing steps differs for different inputs. We accomplish this by using a gating module \mathcal{G}_{ϕ} with parameters ϕ , which, given an input $\mathbf{v}_{i-1} \in \mathbb{R}^{C \times L}$, decides if processing by block \mathcal{B}_{θ_i} is needed (i.e., $\mathcal{G}_{\phi}(\mathbf{v}_{i-1}) = 1$) or not (i.e., $\mathcal{G}_{\phi}(\mathbf{v}_{i-1}) = 0$). Formally, a step of refinement then becomes:

$$F_i(\mathbf{v}_{i-1}; \mathcal{B}_{\theta_i}, \mathcal{G}_{\phi}) = \mathcal{B}_{\theta_i}(\mathbf{v}_{i-1}) \mathcal{G}_{\phi}(\mathbf{v}_{i-1}) + \mathbf{v}_{i-1} (1 - \mathcal{G}_{\phi}(\mathbf{v}_{i-1})). \quad (7)$$

Thus, the separation module \mathcal{S} can then be formulated as follows:

$$\mathcal{S}(\mathbf{v}) = (F_n \circ F_{n-1} \circ \dots \circ F_1)(\mathbf{v}). \quad (8)$$

The gating module used in the experiments follows a similar architecture as described in [20]. Specifically,

$$\mathcal{G}_{\phi}(\mathbf{v}_{i-1}) = \begin{cases} \arg \max (f_{\phi}(\mathbf{v}_{i-1}) + G) & \text{Forward pass} \\ \text{softmax}(f_{\phi}(\mathbf{v}_{i-1}) + G) & \text{Backward pass} \end{cases}, \quad (9)$$

where f turns \mathbf{v}_{i-1} into a two-dimensional vector, G is a two-dimensional vector of Gumbel noise [22], and we abuse the softmax notation to indicate the corresponding probability of being equal to 1. We then define g as the estimated total number of processing steps performed, which is calculated as the sum of all the output values $\mathcal{G}_{\phi}(\mathbf{v}_{j-1})$, $j \in [1, n]$:

$$g = \sum_{j=1}^n \mathcal{G}_{\phi}(\mathbf{v}_{j-1}). \quad (10)$$

The gating module is trained along with the separation network using an auxiliary loss $\mathcal{L}_{\mathcal{G}}$ based on g . This loss is then added to the loss which quantifies separation performance $\mathcal{L}(\hat{\mathbf{s}}, \mathbf{s})$. During inference, Gumbel noise is not used in the gating module, meaning that the choice of the gating module is deterministic. Therefore, we note that

once a skip decision has been taken (i.e., $\mathcal{G}_\phi(\mathbf{v}_{i-1}) = 0$), we can take an early exit as shown in Fig. 4.




Fig. 4: Separation module \mathcal{S} paired with a gating module \mathcal{G}_ϕ which during inference decides at each step of the refinement process whether an early exit should be taken or not.

3. EXPERIMENTAL SETUP

3.1. Dataset

For our experiments, we use the WHAM! dataset [23], which is constructed using two-speaker utterances from the wsj0-2mix dataset [24] and non-stationary ambient noise sources. The dataset consists of 20,000, 3,000, and 5,000 training, validation, and test samples, respectively. The two speaker utterances are mixed with a uniform random input SNR sampled between the range $\mathcal{U}[0, 5]$ dB, while the SNR value between the first (louder) speaker and the noise waveform is similarly sampled from $\mathcal{U}[-3, 6]$ dB.

We perform experiments on two tasks, namely, noisy speech separation and speech enhancement, where each mixture consists of two utterances plus the additive noise and one speech utterance plus the noise waveform, respectively.

3.2. Separation Network

We choose the *Sudo rm -rf* [25] architecture for our experiments, as its performance is close to state-of-the-art architectures while requiring lower computational and memory complexity. The encoder and decoder modules use a kernel size of 21 samples and stride of 10 samples, while the number of bases of the encoder and the decoder is set to 512. Each U-ConvBlock serves as a sub-block and consists of 5 successive downsampling and upsampling operations.

3.3. Training and Evaluation Details

We train all models for 200 epochs on 4-second mixture chunks (zero-padded if needed) sampled at 8 kHz. We use the Adam optimizer [26] with an initial learning rate set to 10^{-3} which decays to one-third of its previous value every 40 epochs. Additionally, we perform gradient clipping when the L_2 -norm exceeds 5. We select a batch size of $B = 4$ and we perform online data augmentation [27] within each batch. All models are trained to minimize the negative scale-invariant signal to distortion ratio (SI-SDR) [28] loss, computed between the estimated source $\hat{\mathbf{s}}$ and the clean source \mathbf{s} which is defined as:

$$\text{SI-SDR}(\hat{\mathbf{s}}, \mathbf{s}) = 10 \log_{10} (\|\rho \mathbf{s}\|^2 / \|\rho \mathbf{s} - \hat{\mathbf{s}}\|^2), \quad (11)$$

where $\rho = \hat{\mathbf{s}}^\top \mathbf{s} / \|\mathbf{s}\|^2$ is a scalar. The loss function used during speech separation training is the permutation invariant [29, 30] negative SI-SDR, where permutations are only considered on the speech sources. For the enhancement task, we use the mean negative SI-SDR between the estimated speech and the ground-truth target.

Table 1: Mean test SI-SDR (dB) performance for block-wise iterative latent refinement for the tasks of speech enhancement (Enh.) and separation (Sep.) where each block consists of several sub-blocks and might be used for a number of iterations. Non-iterative models are denoted with an \times in the iterations columns.

Model Configuration			SI-SDRi (dB)		Params (M)		
Blocks	Sub-Blocks	Iter.	Enh.	Sep.	Enh.	Sep.	
1	1	\times	10.2	7.1			
		2	11.5	8.8			
		3	12.0	9.5	0.4	0.5	
		4	12.3	10.0			
		16	13.6	12.3			
		\times	13.0	10.9			
	4	2	13.6	12.2	0.8	1.0	
		3	14.0	12.9			
		4	14.2	13.6			
		\times	13.8	12.6			
	8	\times	13.8	12.6	1.4	1.5	
		16	\times	14.4	13.8	2.6	2.7
	2	4	2	14.3	13.9	1.4	1.5
	Sepformer [31]			14.4	16.3	25.7	25.7

During model evaluation, we report the mean SI-SDR improvement (SI-SDRi) of the speech sources, which quantifies the improvement the estimated sources $\hat{\mathbf{s}}$ achieve over using the input mixture signal as a source estimate.

4. RESULTS & DISCUSSION

4.1. Iterative Block-Wise Latent Refinement

In the first set of experiments, we compare the performance and model size of the block-wise iterative method (as presented in Section 2.1) against conventional non-iterative models with various numbers of blocks, sub-blocks, and iterations. Each single block can be conveniently defined using several U-ConvBlocks (sub-blocks) from the *Sudo rm -rf* architecture. For example, a conventional *Sudo rm -rf* network with k U-ConvBlocks (each with unique parameters) can alternatively be viewed as a single block containing k sub-blocks and without iterations.

In Table 1, we summarize the results for our iterative block-wise latent refinement method versus conventional non-iterative models under both speech separation and speech enhancement setups. Looking at the results with 1 block, more iterations lead to significantly higher SI-SDR performance without increasing the trainable parameters of the overall model: $10.2 \rightarrow 13.6$ dB and $7.1 \rightarrow 12.3$ dB for separation and denoising, respectively. Overall, all models that reuse their blocks lead to results similar to those of the conventional models while requiring a much smaller model size, making them more amenable to deployment on edge devices. Moreover, when increasing the representational capability of the block using four sub-blocks, we obtain a separation performance close to the state-of-the-art Sepformer [31], which contains $30\times$ more parameters and requires significantly more memory and computational resources. The latter result highlights the importance of employing sophisticated iterative latent refinement strategies to better understand and reduce the redundant sub-structure of source separation networks.

4.2. Progressive Latent Representation Training

In the next set of experiments, we test the proposed progressive training scheme (as presented in Section 2.2). Our aim is to compare it against end-to-end training in terms of performance and memory requirements during training. An additional goal is to investigate the effect of the number of training steps. We note that the number of progressive training steps coincides with the number of blocks. The blocks are in some experiments iterated and contain a number of sub-blocks which again correspond to U-ConvBlocks. We split the total 200 training epochs evenly among the progressive training steps, and the learning rate schedule is reset at the start of each step.

Table 2: Progressive Training Results. The number of blocks corresponds to the number of progressive training steps and each of those blocks contains a set number of sub-blocks. The number of iterations refers to the number of times each block is iterated. The right-most columns indicate the maximum memory required in GBs during a backward pass with a batch size of 1.

Model Configuration			SI-SDRi (dB)		GPU Mem (GB)	
Blocks	Sub-Blocks	Iter.	Enh.	Sep.	Enh.	Sep.
2	8	✗	14.0	13.8	1.06	1.09
	4	2	14.0	13.5	1.05	1.08
	2	4	13.8	13.3	1.05	1.07
4	4	✗	13.6	12.9	0.58	0.61
	2	2	13.4	12.7	0.58	0.60
	1	4	13.2	12.2	0.57	0.60
1	16	✗	14.4	13.8	2.02	2.05

We show the separation performance and the actual memory footprint of the proposed progressive training versus the conventional end-to-end training in Table 2. Progressive training leads to marked improvements in terms of lowering the overall training footprint of separation models which unequivocally enables much higher data-processing parallelization on the training device by allowing the usage of much larger batch sizes. Specifically, when progressively training two blocks with the application of various iterations of each block, we are able to obtain similar separation and denoising performance while halving the total training memory footprint. By increasing the number of progressively trainable blocks to 4 but decreasing the total amount of sub-blocks and iterations in order to keep a constant inference FLOP count, we obtain a remarkable decrease in the overall required training memory bottleneck while observing only a slight performance hit in terms of absolute SI-SDR.

4.3. Adaptive Early Exit

In the last experiment, we propose to use a gating module at the input of each block which dynamically adapts the computation performed for each input mixture. In essence, this adaptive inference scheme decides the number of processing steps performed.

The gating module first transforms the latent representation \mathbf{v}_i of size 128×3200 into a two-dimensional vector using the function f , and then into a binary decision according to Eq. 7. For the function f , we choose a sequence of a Conv1d layer with a kernel of size 1 with 128 input and 2 output channels, a PReLU activation and another Conv1d layer with kernel size of 3200 and 2 input and 2 output channels. The total number of parameters of the gating module is 13K, whereas a single U-ConvBlock contains 147K parameters. The FLOPS overhead is also negligible, namely, a 0.2% increase.

The network that we attach the gating module onto consists of one block containing one U-ConvBlock iterated up to 4 times. This corresponds to the description of Section 2.3 with all blocks having shared parameters. The network without the gating module is pre-trained for 180 epochs. Then both the network and the gating module are jointly trained for 20 more epochs. The initial learning rate is set to 10^{-4} and decays every five epochs to one-third of its previous value. Finally, the overall loss function penalizes bad reconstruction quality as well as enforces the model to use less iterations:

$$\mathcal{L} = \mathcal{L}_{\text{SI-SDR}} + \mathcal{L}_{\mathcal{G}}, \quad \mathcal{L}_{\mathcal{G}} = 0.75(g - 3)^2, \quad (12)$$

where g is the estimated number of iterations performed (see Eq. 10).




Fig. 5: Comparison of the adaptive model (grey) against the baseline models (blue and gold) that have been trained to perform 3 and 4 iterations respectively. Scatter plot of mean SI-SDR performance and mean number of iterations of samples grouped by their SNR. Points are annotated with the corresponding quartile bin. For example, points annotated with 1 represent a group of samples with SNR in the bottom 25%, while points annotated with 4 represent the top 25%. Each marker color/shape represents a different model.

In Fig. 5, we show the behavior of the model with the adaptive gating module. We observe that the model manages to distinguish samples based on their input SNR and adapt to them, while performing 3.6 iterations on average. Specifically, the gating module chooses to perform on average more (less) iterations for samples with lower (higher) input SNR, thus, saving computational resources and potentially being able to better fit the data. Moreover, we compare the adaptive model against two baseline models with a) 1 block with 3 iterations and b) 1 block with 4 iterations. First, we note that in terms of mean SI-SDR and in most of the quartile bins, the adaptive model outperforms the baseline with 3 iterations. Compared to the baseline with 4 iterations, the adaptive model leads to a small separation performance drop but this can be explained due to the fact that the adaptive model has to maintain high-level performance for both 3 and 4 iterations, while the baseline model is allowed to focus only on performing well for 4 iterations.

5. CONCLUSION

We have provided a general formulation for source separation model training as a latent variable refinement process. Our holistic approach enables several optimizations and improvements over conventional models and end-to-end training by explicitly leveraging the repetitive structure of processing blocks inside a neural network. Our experiments showcase the effectiveness of our approach in terms of providing a) a much lighter computational model by leveraging the application of processing blocks multiple times, b) a much lower training memory footprint by using progressive training, and c) a much more flexible inference graph which can dynamically adapt its structure given the input data.

6. REFERENCES

- [1] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle, “Greedy layer-wise training of deep networks,” *Proc. NeurIPS*, 2006.
- [2] T. Karras, T. Aila, S. Laine, and J. Lehtinen, “Progressive growing of GANs for improved quality, stability, and variation,” in *Proc. ICLR*, 2018.
- [3] L. Q. Trinh, “Greedy layerwise training of convolutional neural networks,” Ph.D. dissertation, Massachusetts Institute of Technology, 2019.
- [4] T. Gao, J. Du, L.-R. Dai, and C.-H. Lee, “Densely connected progressive learning for lstm-based speech enhancement,” in *Proc. ICASSP*, 2018, pp. 5054–5058.
- [5] J. Llobart, D. Ribas, A. Miguel, L. Vicente, A. Ortega, and E. Lleida, “Progressive speech enhancement with residual connections,” in *Proc. Interspeech*, 2019, pp. 3193–3197.
- [6] L. Sun, J. Du, X. Zhang, T. Gao, X. Fang, and C.-H. Lee, “Progressive multi-target network based speech enhancement with snr-preselection for robust speaker diarization,” in *Proc. ICASSP*, 2020, pp. 7099–7103.
- [7] S. Kim and M. Kim, “Bloom-net: Blockwise optimization for masking networks toward scalable and efficient speech enhancement,” in *Proc. ICASSP*, 2022, pp. 366–370.
- [8] R. Socher, C. C. Lin, C. Manning, and A. Y. Ng, “Parsing natural scenes and natural language with recursive neural networks,” in *Proc. ICML*, 2011, pp. 129–136.
- [9] D. Eigen, J. Rolfe, R. Fergus, and Y. LeCun, “Understanding deep architectures using a recursive convolutional network,” *arXiv preprint arXiv:1312.1847*, 2013.
- [10] M. Liang and X. Hu, “Recurrent convolutional neural network for object recognition,” in *Proc. CVPR*, 2015, pp. 3367–3375.
- [11] T. Komatsu, “Non-autoregressive asr with self-conditioned folded encoders,” in *Proc. ICASSP*, 2022, pp. 7427–7431.
- [12] A. Li, C. Zheng, L. Cheng, R. Peng, and X. Li, “A time-domain monaural speech enhancement with feedback learning,” in *Proc. APSIPA ASC*, 2020, pp. 769–774.
- [13] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” in *Proc. NeurIPS*, 2020, pp. 6840–6851.
- [14] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, “Wavegrad: Estimating gradients for waveform generation,” in *Proc. ICLR*, 2021.
- [15] Y. Han, G. Huang, S. Song, L. Yang, H. Wang, and Y. Wang, “Dynamic neural networks: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 7436–7456, 2021.
- [16] Y. Kaya, S. Hong, and T. Dumitras, “Shallow-deep networks: Understanding and mitigating network overthinking,” in *Proc. ICML*, 2019, pp. 3301–3310.
- [17] S. Chen, Y. Wu, Z. Chen, T. Yoshioka, S. Liu, J. Li, and X. Yu, “Don’t shoot butterfly with rifles: Multi-channel continuous speech separation with early exit transformer,” in *Proc. ICASSP*, 2021, pp. 6139–6143.
- [18] A. Li, C. Zheng, L. Zhang, and X. Li, “Learning to inference with early exit in the progressive speech enhancement,” in *Proc. EUSIPCO*, 2021, pp. 466–470.
- [19] X. Wang, F. Yu, Z.-Y. Dou, T. Darrell, and J. E. Gonzalez, “Skipnet: Learning dynamic routing in convolutional networks,” in *Proc. ECCV*, 2018, pp. 409–424.
- [20] Q. Guo, Z. Yu, Y. Wu, D. Liang, H. Qin, and J. Yan, “Dynamic recursive neural network,” in *Proc. CVPR*, 2019, pp. 5147–5156.
- [21] Y. Bengio, “Estimating or propagating gradients through stochastic neurons,” *arXiv preprint arXiv:1305.2982*, 2013.
- [22] E. Jang, S. Gu, and B. Poole, “Categorical reparameterization with gumbel-softmax,” in *Proc. ICLR*, 2017.
- [23] G. Wichern, J. Antognini, M. Flynn, L. R. Zhu, E. McQuinn, D. Crow, E. Manilow, and J. Le Roux, “WHAM!: Extending speech separation to noisy environments,” in *Proc. Interspeech*, 2019, pp. 1368–1372.
- [24] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *Proc. ICASSP*, 2016, pp. 31–35.
- [25] E. Tzinis, Z. Wang, X. Jiang, and P. Smaragdis, “Compute and memory efficient universal sound source separation,” *Journal of Signal Processing Systems*, vol. 94, no. 2, pp. 245–259, 2022.
- [26] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. ICLR*, 2015.
- [27] E. Tzinis, S. Venkataramani, Z. Wang, C. Subakan, and P. Smaragdis, “Two-step sound source separation: Training on learned latent targets,” in *Proc. ICASSP*, 2020, pp. 31–35.
- [28] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “SDR—half-baked or well done?” in *Proc. ICASSP*, 2019, pp. 626–630.
- [29] Y. Isik, J. Le Roux, Z. Chen, S. Watanabe, and J. R. Hershey, “Single-channel multi-speaker separation using deep clustering,” in *Proc. Interspeech*, 2016, pp. 545–549.
- [30] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, “Permutation invariant training of deep models for speaker-independent multi-talker speech separation,” in *Proc. ICASSP*, 2017, pp. 241–245.
- [31] C. Subakan, M. Ravanelli, S. Cornell, M. Bronzi, and J. Zhong, “Attention is all you need in speech separation,” in *Proc. ICASSP*, 2021, pp. 21–25.