

LightSAFT: Lightweight Latent Source Aware Source Frequency Transform for Source Separation

Yeong-Seok Jeong^{*1}, Jinsung Kim^{†1}, Woosung Choi², Jaehwa Chung³,
and Soonyoung Jung^{‡1}

Arxiv DOI: [10.21105/joss.01667](https://doi.org/10.21105/joss.01667)

1 Korea University 2 Queen Mary University of London 3 Korea National Open University

Abstract

Conditioned source separations have attracted significant attention because of their flexibility, applicability and extensionality. Their performance was usually inferior to the existing approaches, such as the single source separation model. However, a recently proposed method called LaSAFT-Net has shown that conditioned models can show comparable performance against existing single-source separation models. This paper presents LightSAFT-Net, a lightweight version of LaSAFT-Net. As a baseline, it provided a sufficient SDR performance for comparison during the Music Demixing Challenge at ISMIR 2021. This paper also enhances the existing LightSAFT-Net by replacing the LightSAFT blocks in the encoder with TFC-TDF blocks. Our enhanced LightSAFT-Net outperforms the previous one with fewer parameters.

License

Authors of papers retain copyright and release the work under a Creative Commons Attribution 4.0 International License ([CC BY 4.0](https://creativecommons.org/licenses/by/4.0/)).

In partnership with



Introduction

Recently, many methods based on machine learning have been conducted for music source separation. They can be distinguished depending upon strategies for separation as follows.: single source separation (Choi et al., 2019; Jansson et al., 2017; Takahashi & Mitsufuji, 2017), multi-head source separation (Défossez et al., 2019; Doire & Okubadejo, 2019; Kadandale et al., 2020; Manilow, Seetharaman, et al., 2020; Sawata et al., 2021), one-hot conditioned separation (Choi, Kim, Chung, et al., 2021; Meseguer-Brocal & Peeters, 2019; Samuel et al., 2020; Slizovskaia et al., 2021), Query-by-Example (QBE) separation (Lee et al., 2019; Lin et al., 2021), recursive separation (Takahashi et al., 2019; Wichern et al., 2019), embedding space separation Luo et al. (2017) and hierarchical separation (Manilow, Wichern, et al., 2020). The conditioned source separation method separates target source from a given mixture depending on its conditioning input, which can be various forms (e.g., one-hot conditions, QBE, etc.). Although it has to separate with more complex mechanisms than the multi-head separation method, it is still promising because of its flexibility, applicability and extensionality. For example, it's easy to extend another domain such as text-conditioned audio manipulation (Choi, Kim, Ramírez, et al., 2021). In this paper, we focus on the one-hot conditioned source separation method since it has extensionality.

LaSAFT-Net (Choi, Kim, Chung, et al., 2021) is one of the representative conditioned source separation models. It applies the Latent Source aware Frequency Transformation (LaSAFT) blocks. A LaSAFT block aims to capture the latent source's frequency patterns depending upon a given symbol that specifies which target source we want to separate. Assuming that each latent source contains independent information depending on its viewpoint, it employs the attention mechanism to model the relevance between latent sources and the target symbol. With this approach, LaSAFT-Net showed comparable performance to single-source separation

* co-first author

† co-first author

‡ corresponding author

models on the MusDB18 (Rafii et al., 2017) benchmark even if conditioned models are usually inferior to single-source separation models.

However, it consumes numerous parameters to analyze latent sources. It prevents the model's applicability to a real-world environment, where resources are restricted, such as Music Demix Challenge (MDX) (Mitsufuji et al., 2021) at ISMIR 2021. To this end, we explore a lightweight version of LaSAFT-Net to make it affordable in a real-world environment by decreasing the number of parameters and simplifying the structure.

In this paper, we explore the method for the light version of LaSAFT-Net. The existing LaSAFT-Net uses numerous parameters to separate each latent source. We focus on reducing parameters by introducing shared components in latent source analysis. Furthermore, we found that replacing each LaSAFT block in the encoder with a simple block can improve quality with fewer parameters. We use the Time-Frequency Convolution followed by a Time Distributed Fully-connected networks (TFC-TDF) (Choi et al., 2019) instead of each LaSAFT block.

This paper's contributions are summarized as follows:

- We propose LightSAFT-Net, which is affordable in a restricted environment and shows sufficient performance in MDX challenge.
- We propose an enhanced version of LightSAFT, which shows a better performance than the existing version.

Related work

TFC-TDF block

A TFC-TDF (Choi et al., 2019) block consists of a Time-Frequency Convolution (TFC) (Takahashi & Mitsufuji, 2017) block and a Time Distributed Fully connected layer (TDF) block. A TFC comprises densely connected convolutional blocks containing convolutional neural network layers, Batch Normalization (BN) and ReLU. It extracts high-level features from the localized features of the time-frequency dimension. A TDF was proposed for Frequency Transformation (FT) (Choi, Kim, Chung, et al., 2021; Choi, Kim, Ramírez, et al., 2021; Yin et al., 2020), which aims to capture frequency-to-frequency dependencies of the target source. Choi et al. (2019) showed that replacing fully convolutional intermediate blocks with TFC-TDF blocks improves separation quality. We apply TFC-TDF blocks in the encoder of the original LaSAFT-Net to catch the conditioned source's patterns in an input mixture.

Latent Source

Music source separation aims to separate symbolic-labeled target instruments sets defined by the training dataset. However, some musical sub-components can be grouped as a single instrument even if they have different sonic characteristics. For example, bass-drum and Hi-hat are included in 'drums set' though they have different timbre. Arbitrary instruments set, which do not share the common acoustic features, might make a model hinder separating it from a mixture. It is even more challenging for conditioned source separation models, which modulate internal features according to a given conditioning symbol.

Some existing studies adopted latent source analysis (Choi, Kim, Chung, et al., 2021; Choi, Kim, Ramírez, et al., 2021; Wisdom et al., 2020) to relax the problems caused by symbolically grouped instruments' sets. They assume that a model can learn a latent source that can deal with a more detailed aspect of sound from a given training dataset. For example, a LaSAFT block aims to analyze frequency patterns of latent sources. It aggregates the analyzed patterns by computing an attention matrix between latent sources and a given input conditioning symbol (e.g., one-hot encoding vector that specifies which source a user wants to separate).

In this paper, we apply the concepts of latent source to consider sub-components of grouped instrument's set.

Light-weight latent Source Attentive Frequency Transformation Network

Latent Source Attentive Separation mechanism

We present how the original LaSAFT block processes an input feature as follows. A LaSAFT block analyzes frequency patterns of latent sources using TDF blocks and aggregates the analyzed patterns by computing an attention matrix between latent sources and a given input conditioning symbol. An attention score measures how important the relationship between latent sources and the target is for separation. To compute an attention matrix, it first generates a query vector ($Q \in \mathbb{R}^{1 \times d_k}$), which has d_k dimensions, using an embedding layer. Key vectors ($K \in \mathbb{R}^{|S_L| \times d_k}$) are representative vectors of latent sources, which are learnable parameters. To compute the attention scores, we apply a dot product of the query and key and take softmax after scaled by $\sqrt{d_k}$.

The multiple TDF blocks generate latent source features $V \in \mathbb{R}^{T \times d_k \times |S_L|}$, where the T is the number of frames. The attention mechanism which mixes the latent sources' features is as follows:

$$Attention(Q, K, V) = softmax(QK^T / \sqrt{d_k})V$$

Lightweight latent Source Attentive Frequency Transformation Block

The TDF of the original LaSAFT block contains two phases (Figure 1.(a)). In the first phase, fully connected layers extract latent sources' features from mixture sources. In the second phase, other linear layers are applied to generate the final output. Even if each fully connected layers downsample the internal space, this architecture still consumes many parameters. It seems that it contains too many connections between two phases because a connection exists for each pair of different latent sources. Although such connections might be helpful for separation enabling active communication between sub-modules, it may be computationally difficult to use this model in real-time environment. Therefore, we focus on this point and explore the methods for lightening the model's parameters and maintaining the performance.




Figure 1: Figure 1. Comparison the original LaSAFT block and proposed LightSAFT block. The x is the input intermediate feature and the v_i is generated latent source.

Figure 1 shows the difference between the original LaSAFT block and the proposed LightSAFT block in the latent source separating process. The blocks receive the intermediate feature x and generate the latent source V . Each FC block comprises a fully connected layer (FC), Batch Norm, and activation function in each block. We eliminate inter-source connections between

two phases and use a shared linear layer in the in phase 2. The LightSAFT block has only one separating process in comparison with the original block, simply assuming that the redundant separating process in the original block is not necessary. From this approach, we improve the network’s applicability and efficiency. To validate that our block is reasonable for the MDX challenge, which did not allow GPU computing to infer, we evaluated the performance in its environment.

Advanced LightSAFT

To develop a lightweight LaSAFT-Net, we replace each LaSAFT block with our LightSAFT block. We call the revised separation model LightSAFT-Net for the rest of this paper. The LightSAFT-Net shows sufficient performance with fewer parameters in the MDX challenge.

We further replace each LightSAFT block with a TFC-TDF block in the encoder, to reduce the number of parameters. A TFC-TDF block has fewer parameters than the LightSAFT block because it does not have to generate $|S_L|$ latent source. We call this advanced one LightSAFT-Net+. Although LightSAFT-Net+ has fewer parameters, it performs better than the LightSAFT-Net posted as the MDX challenge comparison.

Experiments

Model Configuration

To fairly compare the number of parameters and performance, we follow the configurations of Choi, Kim, Chung, et al. (2021). The $|S_L|$ is 16 and d_k is 24. Those are the same as the original LaSAFT-Net’s configuration (Table 1). It is less than half of the number of the LaSAFT-Net’s parameters. We trained our proposed model by minimizing the mean squared error between the target’s short-Time Fourier Transformation (STFT) and the model’s output. For data augmentation, we generated the mixtures by mixing the different track’s sources.

Results

model	Runnable?	# of parameters	vocals	drums	bass	other	Avg
LaSAFT-Net	yes	4.5M	-	-	-	-	-
LightSAFT-Net	no	3.8M	6.685	5.272	5.498	4.121	5.394
LightSAFT-Net+	no	2M	7.275	5.935	5.823	4.557	5.897

Table 1. A comparison with LaSAFT-Net. The *Runnable?* means that the model is runnable in MDX condition.

model	type	vocals	drums	bass	other	Avg
Demucs48-HQ (Défossez et al., 2019)	Single	6.496	6.509	6.470	4.018	5.873
XUMX (Sawata et al., 2021)	multi-head	6.341	5.807	5.615	3.722	5.372
UMX (Stöter et al., 2019)	Single	5.999	5.504	5.357	3.309	5.042
LightSAFT-Net	conditioned	6.685	5.272	5.498	4.121	5.394
LightSAFT-Net+	conditioned	7.275	5.935	5.823	4.557	5.897

Table 2. A comparison with other source separation models

We compare the performance between the original model and the proposed models in the same condition.

Table 1 shows the results of the model’s SDR (Vincent et al., 2006) score in the MDX challenge

and the number of each model's parameters. In same condition, the original LaSAFT-Net's parameters are 4.5M, while the LightSAFT-Net has 3.8M parameters sufficient compression for the MDX challenge.

The LightSAFT-Net+, which contains the TFC-TDF blocks in the encoder, has only 2M parameters. The LaSAFT-Net, which cannot separate the songs in a limited time, was not reasonable for this challenge.

On the other hand, the LightSAFT-Net, which is posted as a comparison of the MDX challenge, can separate the music source in a limited time and achieve comparable performance. Even though the LightSAFT-Net+ has the fewest parameters in the comparison group, it shows the best performance. It seems to have no conditioning mechanism, which converts the latent space in encoder inducing stationary training.

Usually, the conditioned source separation models, which can separate all kinds of sources, show inferior performance than the single source separation model; the conditioned model can not focus on the specific task since it have to learn generalized weights for conducting all tasks in limited model complexity. Despite performance degradation, the conditioned source separation model is more attractive because of its applicability and efficiency. Table 2 shows whether the model is conditioned or not and its performance. We take other models' performance at the MDX Leaderboard A, which do not allow additional datasets, for precise comparison.

The LightSAFT-Net+ shows competitive performance despite the conditioned source separation model. Even the model shows a better performance than Demucs-HQ, another baseline of the MDX challenge.

Conclusion

We explore the method to reduce the number of the model's parameters and maintain the performance. We show that our method is reasonable for source separation tasks in the performance and the applicability, which can be used in a restricted environment like the MDX challenge.

The LightSAFT-Net+ shows the competitive performance even though it is a conditioned source separation model. For future work, we will study the method to extend it to other tasks, like TFC-TDF in the encoder.

Acknowledgment

This research was supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education(NRF-2021R1A6A3A03046770). This work was also supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT)(No. NRF-2020R1A2C1012624, NRF-2021R1A2C2011452).

Reference

- Choi, W., Kim, M., Chung, J., & Jung, S. (2021). LaSAFT: Latent source attentive frequency transformation for conditioned source separation. *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 171–175.
- Choi, W., Kim, M., Chung, J., Lee, D., & Jung, S. (2019). Investigating deep neural transformations for spectrogram-based musical source separation. *arXiv Preprint arXiv:1912.02591*.
- Choi, W., Kim, M., Ramírez, M. A. M., Chung, J., & Jung, S. (2021). AMSS-net: Audio manipulation on user-specified sources with textual queries. *arXiv Preprint arXiv:2104.13553*.



- Défossez, A., Usunier, N., Bottou, L., & Bach, F. (2019). Music source separation in the waveform domain. *arXiv Preprint arXiv:1911.13254*.
- Doire, C. S., & Okubadejo, O. (2019). Interleaved multitask learning for audio source separation with independent databases. *arXiv Preprint arXiv:1908.05182*.
- Jansson, A., Humphrey, E., Montecchio, N., Bittner, R., Kumar, A., & Weyde, T. (2017). *Singing voice separation with deep u-net convolutional networks*.
- Kadandale, V. S., Montesinos, J. F., Haro, G., & Gómez, E. (2020). *Multi-task u-net for music source separation*.
- Lee, J. H., Choi, H.-S., & Lee, K. (2019). Audio query-based music source separation. In A. Flexer, G. Peeters, J. Urbano, & A. Volk (Eds.), *Proceedings of the 20th international society for music information retrieval conference, ISMIR 2019, delft, the netherlands, november 4-8, 2019* (pp. 878–885). <http://archives.ismir.net/ismir2019/paper/000108.pdf>
- Lee, J. H., Choi, H.-S., & Lee, K. (2019). Audio query-based music source separation. In A. Flexer, G. Peeters, J. Urbano, & A. Volk (Eds.), *Proceedings of the 20th international society for music information retrieval conference, ISMIR 2019, delft, the netherlands, november 4-8, 2019* (pp. 878–885). <http://archives.ismir.net/ismir2019/paper/000108.pdf>
- Lin, L., Kong, Q., Jiang, J., & Xia, G. (2021). A unified model for zero-shot music source separation, transcription and synthesis. *arXiv Preprint arXiv:2108.03456*.
- Luo, Y., Chen, Z., Hershey, J. R., Le Roux, J., & Mesgarani, N. (2017). Deep clustering and conventional networks for music separation: Stronger together. *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 61–65.
- Manilow, E., Seetharaman, P., & Pardo, B. (2020). Simultaneous separation and transcription of mixtures with multiple polyphonic and percussive instruments. *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 771–775.
- Manilow, E., Wichern, G., & Le Roux, J. (2020). Hierarchical musical instrument separation. *International Society for Music Information Retrieval (ISMIR) Conference*, 376–383. ISBN: [978-0-9813537-0-8](https://doi.org/10.1007/978-0-9813537-0-8)
- Meseguer-Brocal, G., & Peeters, G. (2019). CONDITIONED-u-NET: INTRODUCING a CONTROL MECHANISM IN THE u-NET FOR MULTIPLE SOURCE SEPARATIONS. *Proceedings of the 20th International Society for Music Information Retrieval Conference*.
- Meseguer-Brocal, G., & Peeters, G. (2019). CONDITIONED-u-NET: INTRODUCING a CONTROL MECHANISM IN THE u-NET FOR MULTIPLE SOURCE SEPARATIONS. *Proceedings of the 20th International Society for Music Information Retrieval Conference*.
- Mitsufuji, Y., Fabbro, G., Uhlich, S., & Stöter, F.-R. (2021). Music demixing challenge at ISMIR 2021. *arXiv Preprint arXiv:2108.13559*.
- Rafii, Z., Liutkus, A., Stöter, F.-R., Mimilakis, S. I., & Bittner, R. (2017). *MUSDB18-a corpus for music separation*.
- Samuel, D., Ganeshan, A., & Naradowsky, J. (2020). Meta-learning extractors for music source separation. *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 816–820.
- Samuel, D., Ganeshan, A., & Naradowsky, J. (2020). Meta-learning extractors for music source separation. *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 816–820.



- Sawata, R., Uhlich, S., Takahashi, S., & Mitsufuji, Y. (2021). All for one and one for all: Improving music separation by bridging networks. *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 51–55.
- Seetharaman, P., Wichern, G., Venkataramani, S., & Le Roux, J. (2019). Class-conditional embeddings for music source separation. *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 301–305.
- Slizovskaia, O., Haro, G., & Gómez, E. (2021). Conditioned source separation for musical instrument performances. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29, 2083–2095.
- Stöter, F.-R., Uhlich, S., Liutkus, A., & Mitsufuji, Y. (2019). Open-unmix-a reference implementation for music source separation. *Journal of Open Source Software*, 4(41), 1667.
- Takahashi, N., & Mitsufuji, Y. (2017). *Multi-scale multi-band DenseNets for audio source separation*. <http://arxiv.org/abs/1706.09588>
- Takahashi, N., Parthasaarathy, S., Goswami, N., & Mitsufuji, Y. (2019). Recursive speech separation for unknown number of speakers. In G. Kubin & Z. Kacic (Eds.), *Interspeech 2019, 20th annual conference of the international speech communication association, graz, austria, 15-19 september 2019* (pp. 1348–1352). ISCA. <https://doi.org/10.21437/Interspeech.2019-1550>
- Vincent, E., Gribonval, R., & Févotte, C. (2006). Performance measurement in blind audio source separation. *IEEE Transactions on Audio, Speech, and Language Processing*, 14(4), 1462–1469.
- Wichern, G., Antognini, J., Flynn, M., Zhu, L. R., McQuinn, E., Crow, D., Manilow, E., & Roux, J. L. (2019). WHAM!: Extending speech separation to noisy environments. *arXiv Preprint arXiv:1907.01160*.
- Wisdom, S., Tzinis, E., Erdogan, H., Weiss, R. J., Wilson, K., & Hershey, J. R. (2020). Unsupervised sound separation using mixture invariant training. *NeurIPS*. <https://arxiv.org/pdf/2006.12701.pdf>
- Yin, D., Luo, C., Xiong, Z., & Zeng, W. (2020). PHASEN: A phase-and-harmonics-aware speech enhancement network. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34, 9458–9465.