

DEEP LEARNING BASED SOURCE SEPARATION APPLIED TO CHOIR ENSEMBLES

Darius Petermann¹ Pritish Chandna¹ Helena Cuesta¹
Jordi Bonada¹ Emilia Gómez^{2,1}

¹ Music Technology Group, Universitat Pompeu Fabra, Barcelona

² European Commission, Joint Research Centre, Seville

dariusarthur.petermann01@estudiant.upf.edu,

{pritish.chandna, helena.cuesta, jordi.bonada, emilia.gomez}@upf.edu

ABSTRACT

Choral singing is a widely practiced form of ensemble singing wherein a group of people sing simultaneously in polyphonic harmony. The most commonly practiced setting for choir ensembles consists of four parts; Soprano, Alto, Tenor and Bass (SATB), each with its own range of fundamental frequencies (FOs). The task of source separation for this choral setting entails separating the SATB mixture into the constituent parts. Source separation for musical mixtures is well studied and many deep learning based methodologies have been proposed for the same. However, most of the research has been focused on a typical case which consists in separating vocal, percussion and bass sources from a mixture, each of which has a distinct spectral structure. In contrast, the simultaneous and harmonic nature of ensemble singing leads to high structural similarity and overlap between the spectral components of the sources in a choral mixture, making source separation for choirs a harder task than the typical case. This, along with the lack of an appropriate consolidated dataset has led to a dearth of research in the field so far. In this paper we first assess how well some of the recently developed methodologies for musical source separation perform for the case of SATB choirs. We then propose a novel domain-specific adaptation for conditioning the recently proposed U-Net architecture for musical source separation using the fundamental frequency contour of each of the singing groups and demonstrate that our proposed approach surpasses results from domain-agnostic architectures.

1. INTRODUCTION

Choir music is a well-established and long-standing practice involving a body of singers performing together. Such ensembles are usually referred to as choir and may perform with or without instrumental accompaniment. A choir ensemble is usually structured by grouping the voices into

four different sections, each depicting different frequency ranges for the singers; "Soprano" (260 Hz-880 Hz), "Alto" (190 Hz-660 Hz), "Tenor" (145 Hz-440 Hz), and "Bass" (90 Hz-290 Hz) [1]. This type of structural setting is usually referred to as a SATB setting. Although different variants of this structure exist, the SATB is the most well documented, with several conservatories across Europe dedicated to the study and practice of the art form, highlighting its cultural significance. This will be the main focal point of our study.

The segregation of a mixture signal into its components is a well researched branch of signal processing, known as source separation. For polyphonic music recordings, this implies the isolation of the various instruments mixed together to form the whole. With applications such as music remixing, rearrangement, audio restoration, and full source extraction, its potential use in music is of great appeal. While the task remains similar independently of the type of setting involved, the nature of the sources (e.g.: speech, musical instrument, singing voice) and their relations may entail various challenges and, consequently, require different separation methodologies to be employed.

The most studied case of musical source separation focuses on pop/rock songs, which typically have three common sources; vocals, drums, bass along with other instrumental sources which are usually grouped together as others. A large body of research [2–4] has been published in this field over the last few years, beginning with the consolidation of a common dataset for researchers to train and evaluate their models on. In 2016, *DSD100* [5] was first introduced and made available to the public and was later extended to *MUSDB18* [6], which comprises 150 full-length music tracks for a total of approximately 10 hours of music. To this day, *MUSDB18* represents the largest freely available dataset of its kind.

While source separation for the pop/rock case has come leaps and bounds in the last few years, it remains largely unexplored for the SATB choir case, despite its cultural importance. This is partly due to the lack of a consolidated dataset, similar to the *MUSDB18*, and partly due to the nature of the task itself. The sources to be separated in pop/rock have distinct spectral structure; the voice is a harmonic instrument and has a distinct spectral shape, defined by a fundamental frequency and its harmonic partials



and formants. The bass element to be separated also has a harmonic structure, but lacks the formants found in the human voice and has a much lower fundamental frequency than the human voice. In contrast, the spectrum of a percussive instrument is generally inharmonic and energy is usually spread across the spectrum. In contrast, the sources to be separated in a SATB choir all have a similar spectral structure with a fundamental frequency, partials and formants. This makes the task more challenging than its more studied counter part. However, the distinct ranges of fundamental frequencies in the sources to be separated can be used to distinguish between them, a key aspect that we aim to explore in our study.

We build on top of some recently proposed Deep Neural Network (DNN) models to separate SATB monoaural recordings into each of their respective singing groups and then propose a specific adaptation to one of the models. The rest of the paper is organized as follow: Section 2 presents and investigates some of the recently proposed high performance deep learning based algorithms used for common musical source separation tasks, such as the U-Net [7] architecture and its waveform-based adaptation, Wave-U-Net [8]. Section 3 goes over the dataset curation carried out for this experiment. Section 4 presents our adaptation of the conditioned U-Net model described in [9], with a control mechanism conditioned on the input sources’ fundamental frequency (F0). Section 5 defines the evaluation metrics and methodology used in this experiment. In Section 5.2 we evaluate and compare how existing models and our proposed adaptation perform on the task of source separation for SATB recordings. We then present and discuss the results. Section 6 finally concludes with a discussion around our experiment and provide comments on future research that we intend to carry out.

2. RELATED WORK

While source separation has remained relatively unexplored for the case of SATB choirs, a number of architectures have been proposed over the last few years for musical source separation in the pop/rock case. A comprehensive overview of all proposed models is beyond the scope of this study, but we provide a summary of some of the most pertinent models that we believe can easily be adapted to the case in study.

2.1 U-Net

The U-Net architecture [7], which was specifically developed to process and segment biomedical images, inspired many subsequent audio-related adaptations due to its unprecedented performance.

The original model includes an encoding path, which reduces the initial input into a latent representation (bottleneck) followed by a decoding path, which expands the channels’ receptive field back into its original shape while concatenating the feature maps from the contracting path by the mean of skip connection layers.

One of the first paper to present a U-Net adaptation towards audio source separation was proposed by Jansson et al. [10], where they propose an architecture which specifically targets vocal separation performed on western commercial music (or pop music). The authors present an architecture directly derived from the original U-Net one, which takes spectrogram representations of the sources as input and aims at predicting a soft-mask for the targeted source (either vocal or instrumental). The predicted mask is then multiplied element-wise with the original mixture spectrogram in order to obtain the predicted isolated source. It is worth mentioning that for each of the given sources, a U-Net instance is trained in order to predict its respective mask. In the case of SATB mixtures, four U-Net instances are necessary in order to predict each of the four singing groups.

2.2 Conditioned-U-Net

Depending on the nature of the separation task, its underlying process can easily lead to scaling issues. The conditioned U-Net (C-U-Net) architecture, described in [9], aims at addressing this limitations by introducing a mechanism controlled by external data which govern a single U-Net instance. C-U-Net does not diverge much from the initial U-Net one; as an alternative to the multiple instances of the model, each of which is specialized in isolating a specific source, C-U-Net proposes the insertion of feature-wise linear modulation (*FiLM*) layers [11], which represents an affine transform defined by two scalars - γ and β , across the architecture. This allows for the application of linear transformations to intermediate feature maps. These specialized layers conserve the shape of the original intermediate feature input while modifying the underlying mapping of the filters themselves.

$$FiLM(x) = \gamma(z)x + \beta(z) \tag{1}$$

In eq. (1), x is the input of the *FiLM* layer, γ and β the parameters that scale and shift x based on an external information, z [9]. γ and β modulates the feature maps according to an input vector z , which describes the source to separate. The condition generator block described in Figure 1 represents a neural network embedding the one-hot encoding input z into the most optimal values to be used by the *FiLM* layer.

2.3 Wave-U-Net

In [8], the authors present a time-domain adaptation of the U-Net architecture, which performs the separation operation on the waveform. As the input is a one-dimensional signal, the feature maps are computed directly from the waveform samples through 1D convolution operations. Because Wave-U-Net takes raw waveforms as input, the initial U-Net model has to be adapted accordingly in order to accommodate for the input’s nature. Consequently, the feature maps along both contracting and expanding paths are computed by the means of single-dimensional convolution layers. Both paths contain twelve convolutional layers, each, for a total of 24 layers. In the down-sampling

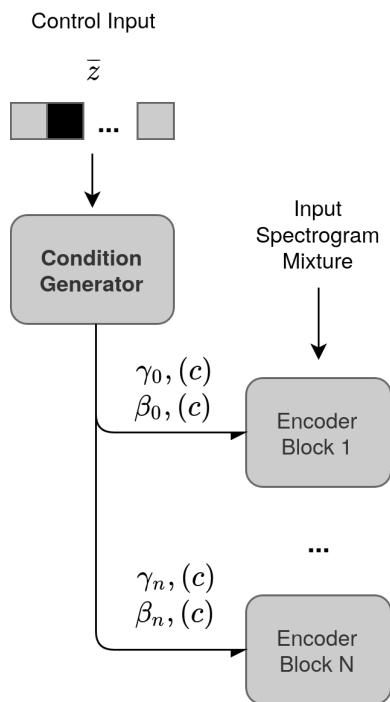


Figure 1: C-U-Net control mechanism with the vector \bar{z} , a one-hot representation of the source to separate, which dictates the N sets of γ and β values to be used by the *FiLM* layer at each of the block of the encoding path, 1 to N.

path, the receptive field is reduced in half after each layer while the input feature maps are increased by a factor of 24 every time. On the other hand, in the up-sampling path the time-context is doubled after every convolutional layer while the feature maps are reduced, again by 24, after every layers. By that mean, the receptive field and the number of channels of the original input signal will remain preserved at the output stage. Although Wave-U-Net has proven to deliver satisfying results on common musical source separation tasks, the fact remains that waveform-based architectures in general require more data than their spectrogram-based counter parts.

3. DATASET

The training data we have curated for this experiment are composed of the following two datasets:

- Choral Singing Dataset [12] (CSD). Three songs performed by 16 singers from the Anton Bruckner Choir (SATB)¹.
- A proprietary dataset with 26 Spanish SATB songs by 4 singers, one for each part.

There are very few publicly available choir music datasets, thus our choice remains limited. To this day and to the best of our knowledge, there isn't any existing dataset which is specifically suited to our task, thus one

of the subsidiary work of this experiment revolves around curating a proper and complete dataset to train our various models. For our experiment, we take advantage of the CSD [12]. This dataset was recorded in a professional studio and contains individual tracks for each of the 16 singers of a SATB choir, i.e. 4 singers per choir section. It comprises three different choral pieces: *Locus Iste*, written by Anton Bruckner, *Niño Dios d'Amor Herido*, written by Francisco Guerrero, and *El Rossinyol*, a Catalan popular song; all of them were written for 4-part mixed choir. The dataset is very well suited for our experiment as the isolated track for each individual singer will allow us to proceed the same way as in [13], that is to create artificial mixes by combining various stems from different groups together. Using different combinations of all 16 singers, we created 256 SATB quartets for each piece, which represent all possible combinations of singers taking into account the voice type restriction (i.e. exactly one singer per voice is needed).

The second dataset we use is a proprietary one including 26 songs for exactly one singer per part (i.e. 4 stems per song), which is a well-suited format for our task as well. All songs offered as part of this dataset are performed in Spanish and their length revolves around two to three minutes, for a total of 58 minutes of audio data.

Our curation work consists in consolidating these two datasets and make sure all the data that we are using remain consistent and well-formatted across all the audio stems, which includes length and amplitude normalization, as well as properties standardization. Most of the files included as part of the initial datasets were presented as 10-seconds long snippets as opposed to full-length songs, which isn't an ideal format to work with. Hence, some additional efforts have been devoted to turn these files in a more convenient and consolidated format.

4. APPROACH

Injecting domain knowledge in DNNs has been proved to be an effective way to learn complex input-output relations with high accuracy when the available data happen to be scarce and limited [14], such as found in our case.

Each of the singing groups in a SATB recordings performs within its own respective frequency range, that is, the voices' F0 contour will rarely overlap across the various groups. This factor makes the sources' F0 a suitable discriminative feature, which could be injected in the DNN during the training stage and potentially improve the separation of the various singing groups in SATB recordings.

In this view, we propose to adapt the original C-U-Net architecture, which initially embeds the instruments to be separated, \bar{z} , in order to produce the various *FiLM* parameters (γ, β), and substitute the external control input data for the F0 track of the target source. The new control input vector \bar{z} will thus hold time as well as frequency dimensions.

¹ <https://zenodo.org/record/1286570#.XyGcHy-z3yU>

4.1 Control Input Representation

As previously mentioned, we use the frame-wise F0 as the external control data of our condition generator network. This entails that a few preliminary steps are required prior to proceeding to the training stage. We first automatically extract the sources' F0 track using the DIO algorithm [15]. Once the raw pitch tracks are obtained, we convert each time-step F0 into a one-hot encoded representation as postulated in [16], that is 60 frequency bins over 6 musical octaves, with a base frequency of 32.7 Hz Hertz, for a total of 360 frequency bins per time-step. As a result, for 128 time-steps, our control input will be in the shape of [128, 360].

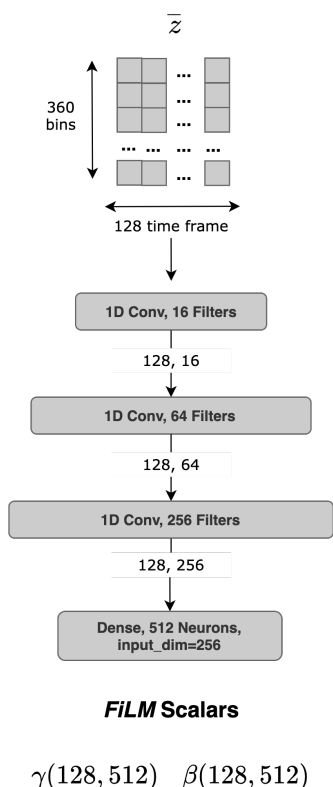


Figure 2: Control model architecture. The convolution is performed across the frequency bins for each time-step. The dense layer provides a specific conditioning for each frequency bin.

4.2 Control Model

The control model used in our proposed architecture embeds the one-hot encoded CQT F0 representation for a given time-step into a set of transforms of identical shape as the spectrogram input. This is achieved by modeling the condition vectors as 1-D data with multiple feature channels. The condition vectors are then fed into a convolutional neural network (CNN) with a kernel of size 10 seizing contextual information from the adjacent time-steps. As a result, all input channels of the initial convolution contribute to all resulting feature maps in the output of the first convolutional layer. Finally a dense layer provides a

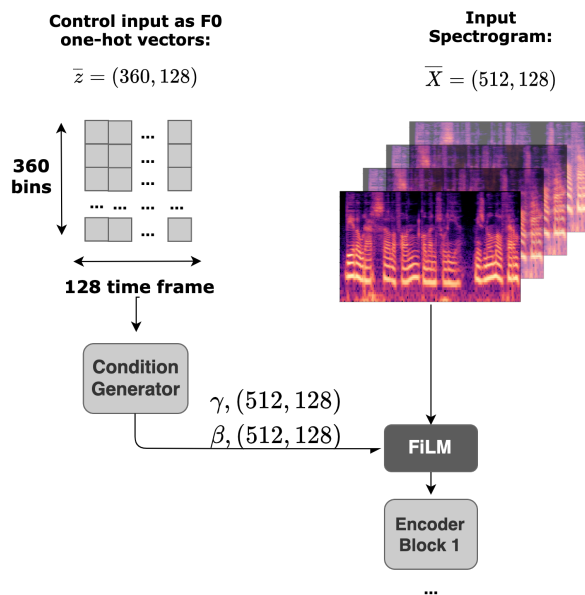


Figure 3: C-U-Net Control Mechanism adapted to our task, with the one-hot vector \bar{z} depicting the various SATB singing groups' F0 contour.

specific conditioning for each frequency bin at each time-step, taking into account the contextual information previously captured by the CNN. Figure 2 shows the condition generator architecture in greater details.

As the temporal relation between the external control input data and the input spectrogram is crucial, it is important to apply these affine transformations while the receptive field of the input is still intact. Hence the *FiLM* layer is applied prior to the encoding path. Figure 3 shows the overall structure behind the proposed conditioning architecture.

We propose two variants of the architecture described above, each of which differs slightly in the way it embeds the control input data; the first variant applies a unique affine transform for each individual frequency bin at every input time-step. The resulting scalars in the output of the external CNN model will thus be in the shape [512, 128] for a given input spectrogram of the same time context. On the other hand, the second variant applies a single transform for all frequency bins at a given time-step, resulting in a set of scalars of shape [1, 128] for 128 spectrogram frames given in the input. We refer to the two approaches as "Domain-Specific Global" and "Domain-Specific Local", respectively and define them with the acronyms *C-U-Net D-S G* and *C-U-Net D-S L* in the rest of this paper. while the three models covered in Section 2 will be referred to as "Domain-Agnostic" models.

5. EVALUATION

To assess our proposed approach and show that injecting domain knowledge as control input data to the network improves its performance on SATB recordings, we evaluate the performances of three domain-agnostic state-of-the-art DNN models; U-Net, its waveform adaptation

Wave-U-Net, and the original C-U-Net. We then compare the results with the two domain-specific models proposed in Subsection 4.2. We evaluate the performance of a model by computing three metrics, SDR, SIR, and SAR [5], between the predicted and true audio sources. The three measurement metrics describe the overall quality of the separation, the level of interference with other sources as well as the amount of artifact added by the separation algorithm, respectively. The metrics are computed using the *mir_eval toolbox* [17] for each of the SATB singing groups.

5.1 Train - Test Split

Given the limited size of our data, we opted to set apart one song from the proprietary dataset as well as one singer per voice for each song from the CSD in order to build our first use-case test set. The rest of the data was used for training. This allowed us evaluate the model on unseen songs and singers. Our second test case contains unison singing, which was not seen at all during training. As such, we used the three songs from the CSD with all singers for evaluation.

5.2 Experiment Results

Subsections 5.2.1 and 5.2.2 present the performance results for the two test cases described earlier; that is, for the test set involving exactly one singer per part and the other involving exactly four singers per part, respectively. *C-U-Net D-A* refers to the domain-agnostic architecture while *C-U-Net D-S L* and *C-U-Net D-S G* refer to our two domain-specific adaptations. For testing, we use the oracle fundamental frequency of each of the sources, pre-computed prior to model inference. In a complete source separation pipeline, we would complement our system with the multi-pitch algorithm proposed by [18].

5.2.1 Use-Case 1:

Table 1 portrays the mean SIR and SAR results on all the SATB parts and average for all five models mentioned in previous sections. Figure 4a details the SDR score distributions on our first test set. We observe that our two adaptations, *C-U-Net D-S L* and *C-U-Net D-S G*, call attention to a significant score gap between domain-agnostic and domain-specific models, with an average increase of about 1dB SDR and 1.5dB SIR between the two different approaches. These improvements underline an overall better quality of the predicted sources (SDR) as well as a decline in interference between the various predictions. Our domain-specific architecture hence demonstrates a better ability to cope with the correlated nature of the various SATB sources and seem to predict an appropriate spectral mask for each of them. We also observe that our proposed adaptations return the lowest SDR, SIR and SAR performances for the Bass part, specifically. This could be due to the fact that the Bass group, among all SATB groups, shares the highest number of harmonics with its other source counter parts.

We also note that the mean SDR substantially drops for the Tenor singing group across nearly all domain-agnostic

models, reaching a negative result with the *C-U-Net D-A* architecture (-1.25dB SDR). We speculate that the reason behind such decline can be directly related to the close nature of the F0 contours of both Alto and Tenor singing groups, making it harder for domain-agnostic architecture to distinguish between the two sources. This limitation brings yet another justification for the conditioning approach we have taken in this paper.

5.2.2 Use-Case 2:

Figure 4b as well as the bottom portion of Table 1 presents the SDR, SIR and SAR scores on the second use-case test set. We observe that the introduction of more complex mixtures involving a higher number of singers (i.e.: 16 singers in this case) decreases the performance of our proposed models, with an average SDR barely surpassing U-Net’s for our *C-U-Net D-S G* model and levelling it out for the *C-U-Net D-S L* model. This can be attributed to the use of the mean of the various pitches present in a singing group, to represent the pitch of the unison. Since domain-agnostic models, such as the plain U-Net, don’t hold this assumption, these architectures are less prone to errors when exposed to these type of mixture settings ².

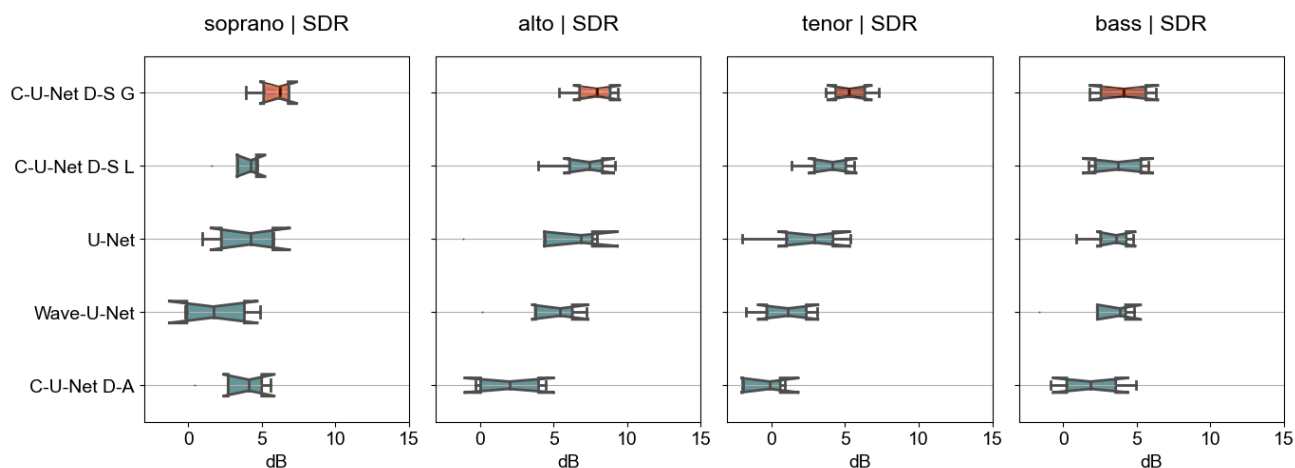
6. CONCLUSIONS AND FUTURE WORK

In this work we have presented the task of musical source separation applied to SATB choir recordings. We first described the consolidated dataset that we’ve specifically curated for this experiment and its potential use and application for future related research. We then assessed how well recent domain-agnostic deep learning based architectures for musical source separation performed on this task, given two different use-cases; 4-singers mixture and 16-singers mixture separation. An adaptation of the U-Net architecture was then proposed, consisting in conditioning some of the network parameters on the fundamental frequency contour of each of the SATB mixture sources. The preliminary results showed that taking advantage of domain-knowledge during the training process improved the performance on both of our proposed use-cases. For the evaluation presented in this paper, we use the oracle F0 is currently used as external control input data to the network. In a complete source separation pipeline, we plan on combining the task of multi-pitch tracking [18] with the system presented in this paper. We also plan on validating our evaluation with perceptual listening tests and exploring applications of the SATB separation. including remixing, transcription and transposition combined with the work presented in [19, 20].

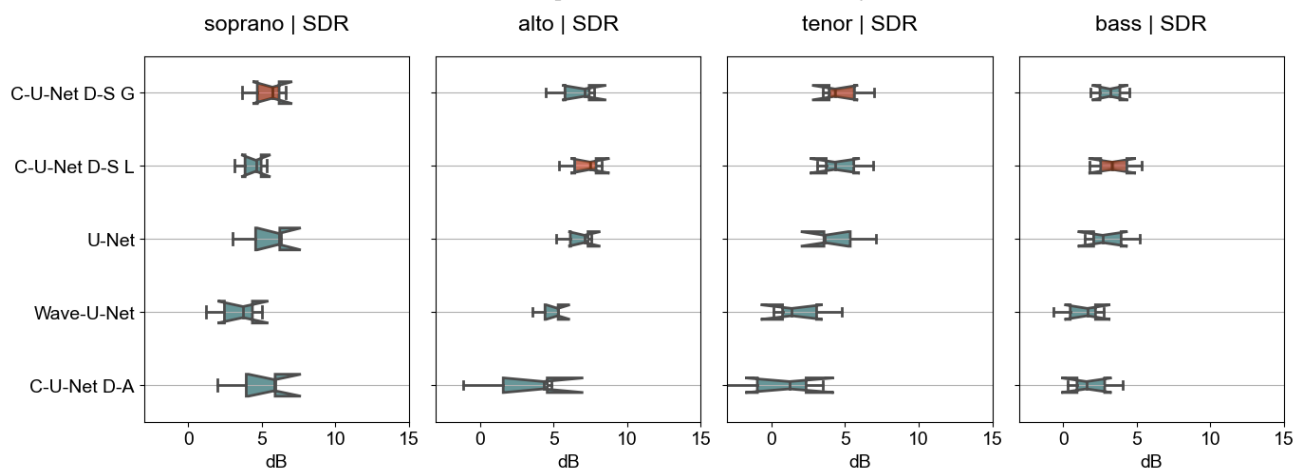
7. ACKNOWLEDGEMENTS

The TITANX used for this research was donated by the NVIDIA Corporation. This work is partially supported by

² Audio examples showcasing how the inference of our proposed models compare against other state-of-the-art architectures are available online. On the same page is also included a link to our pre-trained models: <https://darius522.github.io/satb-source-separation-results/>



(a) Four-Parts SDR Boxplot Results, Use-Case 1: 4-Singers Mixture



(b) Four-Parts SDR Boxplot Results, Use-Case 2: 16-Singers Mixture

Figure 4: Boxplot SDR results on the five U-Net based models described in previous sections. Subfigure 4a shows the result distribution over the first use-case test set while 4b depicts the results for the second use-case. For each one of the SATB parts, the model performing with the highest median is indicated in a dark orange color.

Model	Test Use-Case 1 - SIR (dB)					Test Use-Case 1 - SAR (dB)				
	Soprano	Alto	Tenor	Bass	Avg.	Soprano	Alto	Tenor	Bass	Avg.
<i>Wave-U-Net</i>	5.99±2.4	9.19±2.9	4.62±2.1	8.49±3.5	7.07	5.36±1.7	7.11±2.4	4.79±1.4	4.89±1.4	5.54
<i>U-Net</i>	10.28±2.4	10.77±4.1	6.70±3.2	9.45±2.0	9.30	5.35±1.8	7.13±3.0	5.32±1.8	4.94±1.1	5.69
<i>C-U-Net D-A</i>	10.09±2.6	7.81±1.6	3.32±3.3	7.61±2.2	7.21	5.19±1.6	4.41±2.8	2.65±1.2	4.12±2.0	4.09
<i>C-U-Net D-S L</i>	9.71±1.7	12.37±1.5	9.89±2.2	9.71±1.7	10.42	5.44±1.0	8.75±2.0	5.58±1.3	5.51±1.7	6.32
<i>C-U-Net D-S G</i>	12.72±1.8	14.04±1.5	11.79±1.5	9.78±2.1	12.08	7.02±1.1	9.02±1.6	6.86±1.5	5.93±1.6	7.21

Model	Test Use-Case 2 - SIR (dB)					Test Use-Case 2 - SAR (dB)				
	Soprano	Alto	Tenor	Bass	Avg.	Soprano	Alto	Tenor	Bass	Avg.
<i>Wave-U-Net</i>	8.13±2.1	10.02±0.9	6.80±2.2	7.45±2.0	8.10	5.75±1.0	6.73±1.1	4.79±1.5	3.23±0.9	5.13
<i>U-Net</i>	12.41±1.8	13.11±1.2	10.26±1.1	8.50±2.3	11.07	6.31±1.4	7.97±1.0	6.59±1.8	5.27±1.1	6.54
<i>C-U-Net D-A</i>	11.99±1.8	9.08±2.8	5.65±3.1	7.60±2.0	8.58	5.77±1.7	4.60±2.9	3.39±1.8	4.15±1.2	4.48
<i>C-U-Net D-S L</i>	10.32±1.1	13.06±1.7	10.77±1.5	8.89±2.2	10.76	6.02±0.9	8.59±1.2	6.45±1.7	5.59±1.1	6.66
<i>C-U-Net D-S G</i>	12.08±1.5	13.50±2.5	12.05±1.3	8.91±2.1	11.63	6.68±1.2	7.70±1.3	6.17±1.6	5.17±0.8	6.43

Table 1: SIR and SAR mean and standard deviation results on the four SATB parts as well as their average for the five U-Net based models described in previous sections. The top table depicts the results obtained from the first use-case test set while the bottom one the second use-case test set.

the Towards Richer Online Music Public-domain Archives (TROMPA H2020 770376) project. Helena Cuesta is supported by the FI Predoctoral Grant from AGAUR (Generalitat de Catalunya). The authors would like to thank Rodrigo Schramm and Emmanouil Benetos for sharing their singing voice datasets for this research.

8. REFERENCES

- [1] M. Scirea and J. A. Brown, "Evolving four part harmony using a multiple worlds model," in *Proceedings of the 7th International Joint Conference on Computational Intelligence (IJCCI)*, vol. 1. IEEE, 2015, pp. 220–227.
- [2] P. Chandna, M. Miron, J. Janer, and E. Gómez, "Monoaural Audio Source Separation Using Deep Convolutional Neural Networks," in *International Conference on Latent Variable Analysis and Signal Separation*, 2017, pp. 258–266.
- [3] F.-R. Stöter, S. Uhlich, A. Liutkus, and Y. Mitsufuji, "Open-unmix - a reference implementation for music source separation," *Journal of Open Source Software*, 2019. [Online]. Available: <https://doi.org/10.21105/joss.01667>
- [4] F. Lluís, J. Pons, and X. Serra, "End-to-End Music Source Separation: Is it Possible in the Waveform Domain?" in *Proceedings of Interspeech 2019*, September 2019. [Online]. Available: <http://dx.doi.org/10.21437/interspeech.2019-1177>
- [5] A. Liutkus, F.-R. Stöter, Z. Rafii, D. Kitamura, B. Rivet, N. Ito, N. Ono, and J. Fontcave, "The 2016 signal separation evaluation campaign," in *International Conference on Latent Variable Analysis and Signal Separation*. Springer, 2017, pp. 323–332.
- [6] Z. Rafii, A. Liutkus, F.-R. Stöter, S. I. Mimitakis, and R. Bittner, "The MUSDB18 corpus for music separation," December 2017. [Online]. Available: <https://doi.org/10.5281/zenodo.1117372>
- [7] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proceedings of the International Conference on Medical image computing and Computer-assisted Intervention*. Springer, 2015, pp. 234–241.
- [8] D. Stoller, S. Ewert, and S. Dixon, "Wave-U-Net: A Multi-Scale Neural Network for End-to-End Audio Source Separation," in *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, 2018.
- [9] G. Meseguer-Brocal and G. Peeters, "Conditioned u-net: Introducing a control mechanism in the u-net for multiple source separations." November 2019.
- [10] A. Jansson, E. Humphrey, N. Montecchio, R. Bittner, A. Kumar, and T. Weyde, "Singing Voice Separation with Deep U-Net Convolutional Networks," in *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, 2017.
- [11] E. Perez, F. Strub, H. de Vries, V. Dumoulin, and A. C. Courville, "FiLM: Visual Reasoning with a General Conditioning Layer," in *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, April 2018.
- [12] H. Cuesta, E. Gómez, A. Martorell, and F. Loáiciga, "Analysis of Intonation in Unison Choir Singing," in *Proceedings of the International Conference of Music Perception and Cognition (ICMPC)*, Graz (Austria), July 2018, pp. 125–130.
- [13] H. Cuesta, E. Gómez, and P. Chandna, "A Framework for Multi-f0 Modeling in SATB Choir Recordings," in *Proceedings of the Sound and Music Computing (SMC) Conference*, Málaga (Spain), April 2019.
- [14] M. Silvestri, M. Lombardi, and M. Milano, "Injecting domain knowledge in neural networks: a controlled experiment on a constrained problem," *ArXiv*, vol. abs/2002.10742, 2020.
- [15] M. Morise, H. Kawahara, and H. Katayose, "Fast and reliable f0 estimation method based on the period extraction of vocal fold vibration of singing voice and speech," in *AES 35th International Conference on Audio for Games*, February 2009. [Online]. Available: <http://www.aes.org/e-lib/browse.cfm?elib=15165>
- [16] R. M. Bittner, B. McFee, J. Salamon, P. Li, and J. P. Bello, "Deep Saliency Representations for F0 Estimation in Polyphonic Music," in *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, 2017.
- [17] C. Raffel, B. Mcfee, E. J. Humphrey, O. N. Justin Salamon, D. Liang, and D. P. W. Ellis, "mir_eval: a transparent implementation of common mir metrics," in *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, 2014.
- [18] H. Cuesta, B. McFee, and E. Gómez, "Multiple F0 Estimation in Vocal Ensembles using Convolutional Neural Networks," in *Proceedings of the 21st International Society for Music Information Retrieval Conference (ISMIR)*, Montreal, Canada (Virtual), 2020.
- [19] P. Chandna, H. Cuesta, and E. Gómez, "A Deep Learning Based Analysis-Synthesis Framework For Unison Singing," in *Proceedings of the 21st International Society for Music Information Retrieval Conference (ISMIR)*, Montreal, Canada (Virtual), 2020.
- [20] P. Chandna, M. Blaauw, J. Bonada, and E. Gómez, "Content based singing voice extraction from a musical mixture," in *Proceedings of the 45th IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 2020, pp. 781–785.