Conditioned Source Separation for Musical Instrument Performances

Olga Slizovskaia, Gloria Haro, and Emilia Gómez

Abstract—In music source separation, the number of sources may vary for each piece and some of the sources may belong to the same family of instruments, thus sharing timbral characteristics and making the sources more correlated. This leads to additional challenges in the source separation problem. This paper proposes a source separation method for multiple musical instruments sounding simultaneously and explores how much additional information apart from the audio stream can lift the quality of source separation. We explore conditioning techniques at different levels of a primary source separation network and utilize two extra modalities of data, namely presence or absence of instruments in the mixture, and the corresponding video stream data.

Index Terms—Single Channel Source Separation, Audio-Visual Analysis, Conditioned Neural Networks

1 INTRODUCTION

Humans combine sensory information of different origins to obtain a comprehensive picture of the world. Actually, the information that we perceive with different senses often originates from the same object or event at the physical level. Thus, sound is a vibration which is commonly produced by an object's movement. In the case of musical instruments sounds, we even can see the movements and associate a particular sound with its source [1], [2]. Also, each of those instruments has its own unique visual characteristics such as shape and color which help us recognize them. While experiencing online listening through video streaming, we are often given all this information together with prior knowledge we may have about a piece, instrumentation or an artist. All this helps us with the interpretation of the music we are exposed to, such that, while listening, we can focus on the individual sources of the sound and identify them. Moreover, such phenomenon as synesthesia can also support audio-visual correspondence studies. Thus, some articles report correlations between loudness and visual size/light intensity, as well as musical timbre and arbitrary visual shapes [3].

In this paper, we focus on Single Channel Source Separation (SCSS). This task is usually solved in the audio domain, but in this work we explore the effects of integrating two additional kinds of context data, namely instruments labels and their visual properties.

We work with audio-visual recordings of musical ensembles with several families of instruments that can be commonly found in a symphonic orchestra such as strings, woodwinds and brass instruments, that is to say, mostly chamber music. Source separation with such a setup is

known to be a very challenging task, and has been approached in the past with multi-channel score-informed methods [4] and timbre-informed methods [5]. It is worth noticing that previous studies operated on multi-channel recordings with different instruments present in each channel, as there are no large datasets with ground truth audio available for each target source.

The music source separation problem has several challenges, to mention a few:

- The instruments within a family can be very similar to one another;
- The number of sources in the mixture is unknown in advance;
- There is a considerable amount of overlap in time and frequency between sources;
- Even for instruments which have essentially different timbre, tone color, and different performance techniques, such as clarinet and viola, some musicians may mimic a sound of one while playing another [6].

Different recent works have tackled the music source separation problem using audio-visual data and deep learning techniques [2], [7], [8], [9], [10]. Due to the lack of ground-truth data, all of these works have followed the *mix-and-separate* strategy for training, proposed in [7], which consists of artificially creating mixes by adding the sound of individual sources. Similarly, the same kind of unrealistic mixtures is used for the quantitative evaluation. In addition, previous studies only considered mixes of only two sources, where instruments from the same family were rarely present. These restrictions affect the generality of the proposed approaches. In fact, separating different musical instruments playing the same piece, i.e. a real mixture, is a more challenging task than these unrealistic scenarios, since the different audio sources are synchronized and playing in harmony.

In this work, we also adopt the mix-and-separate technique while training; again, because of the absence of appropriate datasets for our purposes. But contrarily to previous

• O.Slizovskaia and G.Haro are with the Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain, 08018.

E-mail: olga.slizovskaia@upf.edu, gloria.haro@upf.edu

• E.Gómez is with the Joint Research Centre, European Commission, Seville, Spain, 41092; and the Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain, 08018.

Manuscript received March 23, 2020; revised October 13, 2020; accepted April 4, 2021.

works, we evaluate and compare the different methods in a more realistic scenario, thanks to the URMP dataset [11].

As for combining different modalities of information, the key technical problem for many years has been the large gap (both in dimensions and content) between representations of the auditory and visual modalities [12]. One of the common approaches consisted in feature construction followed by dimensionality reduction [12], [13]. Deep learning techniques and their powerful representation capability have helped in reducing the heterogeneity and dimensionality gap among different modalities [14]. However, a proper way of fusing different data representations still remains an issue [15], [16], [17].

In this regard, this work explores conditioning techniques at different levels of a primary source separation network. We are not the first ones to propose Conditioned-U-Net for source separation or audio-visual source separation [2], [7], [9], [18], [19]. However, unlike prior approaches, we conduct a thorough study identifying the optimal type of conditioning and comparing possible conditioning strategies with two types of context data: the presence or absence of instruments in the mixture and the video stream data. We train on mixtures of up to 7 sources, and evaluation is carried out on real-world mixtures from the URMP [11] dataset which has up to 4 different instruments per piece, often from the same family. We should note that we only address the separation of sources of different type (in our case different musical instruments) and thus, two or more instances of the same instrument are treated as a single source.

The complexity of the task allows for the present approach to be used as a baseline for future research. In order to facilitate that, the present study is reproducible as we provide pretrained models, code, data and all the training parameters. The supplementary materials and examples are available at <https://veleslavia.github.io/conditioned-u-net/>.

This paper provides an overview of existing techniques for source separation, audio-visual methods and conditioning strategies in Section 2, indicating relations and differences with the present work. In Section 3 we formalize our approach for performing source separation conditioned on context data. This is followed by Section 4 where we describe the experimental setup and implementation details. Finally, we discuss the obtained results in Section 5 and provide conclusions in Section 6.

2 RELATED WORK

2.1 Single channel source separation

Single channel source separation (SCSS) consists of estimating the individual sources x_i given a mono time-domain mixture signal y of N sources:

$$y(t) = \sum_{i=1}^N x_i(t). \quad (1)$$

Instead of predicting time-domain signals, a general approach for solving SCSS involves the estimation of N masks for Short-Term Fourier transform (STFT) values of the mixture. In this case, we consider a time-frequency

representation of the mixture $\mathbf{Y}$ and of the sources $\mathbf{X}_i$, and the goal of the source separation method is to learn a real-valued (or complex-valued) mask M_i for each source i .

Let us denote the STFT magnitude of $\mathbf{X}_i(\tau, \omega)$ and $\mathbf{Y}(\tau, \omega)$ by $|\mathbf{X}_i|$ and $|\mathbf{Y}|$, respectively, with frequency ω and time frame τ . In this work, we only consider two types of *real-valued* masks, namely *ideal ratio* or *soft* masks M_i^{ir} :

$$M_i^{ir}(\tau, \omega) = \frac{|\mathbf{X}_i(\tau, \omega)|}{|\mathbf{Y}(\tau, \omega)|}, \quad (2)$$

and *ideal binary* masks¹ M_i^{ib} :

$$M_i^{ib}(\tau, \omega) = \begin{cases} 1, & \text{if } \frac{|\mathbf{X}_i(\tau, \omega)|}{|\mathbf{Y}(\tau, \omega)| - |\mathbf{X}_i(\tau, \omega)|} \geq 1 \\ 0, & \text{otherwise,} \end{cases} \quad (3)$$

We obtain the STFT magnitude values of separated sources by multiplying the STFT magnitude of the mixture by the estimated masks $\hat{M}_i$, i.e., $|\hat{\mathbf{X}}_i| = \hat{M}_i \odot |\mathbf{Y}|$, where $\odot$ denotes the Hadamard product (element-wise multiplication). Then, the waveforms of the source signals are recovered by applying the inverse STFT transform on the predicted magnitude $|\hat{\mathbf{X}}_i|$ and using the phase of the mixture $\mathbf{Y}$.

The mask estimation step has always been an essential component of model-based source separation algorithms [4], [5], [20], [21], [22], [23]. Consecutively, the masking-based approach for training neural networks has received a lot of attention recently and has been very successful in SCSS [24], [25], [26]. Additional constraints such as mixture-consistency and STFT consistency on the estimated sources have proven to improve mask-based SCSS networks [25]. Despite the fact that most of the existing work has focused on estimating binary or ratio masks, the direct estimation of STFT magnitude values has also been used in practice [27] together with loss function computation in time-frequency [28] or time domain [29] while internally estimating the masks.

It's worth noting that the set of methods which has been successfully used in source separation is very diverse, and the optimal choice of an architecture remains an open research question. Some examples include LSTMs [30] and BLSTMs [28], [31], fully-connected architectures [32], U-Nets [26], [27], GANs [33], [34], as well as combinations of the above [29], [31]. Some research works suggested the estimation of each source separately with a dedicated network [24], [28], while other approaches employed one-to-many encoder-decoder networks with a shared encoder and one decoder per source [27]. Overall, the use of an individual network for each source seems to provide a better performance but it comes at the cost of increased training time.

There have been diverse proposals for loss functions, which include L_2 -distance [24], [31], and L_1 -distance [26], [27] on estimated spectrograms, L_2 -distance on ratio and binary masks [32], L_1 -distance on ratio masks [9], binary cross entropy on binary masks [2], [7], as well as negative Scale-Invariant Signal to Distortion Ratio (SI-SDR) [30], [35] and Signal to Noise Ratio (SNR) [29] as objective functions.

1. <https://github.com/sigsep/sigsep-mus-oracle/blob/master/IBM.py>

2.2 Audio-visual approaches and source separation

2.2.1 Audio-visual model-based methods

Audio and visual information are related, and we can often see or imagine the source of origin for every particular sound. In the real world, they have a causal relation. In addition, the study of some misattribution effects (i.e. ventriloquism effect, see [13] and references therein) has shown that people tend to relate audio and visual events if they happen simultaneously. Having this in mind, a correlation approach for source localisation was proposed as early as in 2000 [13]. It consisted of calculating intensity changes in audio and video and computing correlations between audio and every pixel in a sequence of frames. The authors showed that the method can successfully identify the speaking person at every time frame in videos of two people speaking in turns.

Thinking along the same line, Kidron et al. presented a method that detects pixels associated with a sound source while filtering out other dynamic pixels [12]. The method used a refined version of canonical correlation analysis and, in contrast to previous studies, which mostly focused on speech applications, it could handle different types of sounding sources, not only people speaking but also musical instruments. The authors also discussed the *chorus ambiguity* phenomenon when several people sing in synchrony, and in this particular case they accepted the detection of any of the faces as a successful result. The main concern raised by the authors is the extreme *locality* of the pixel regions associated with an audio event which they overcome by introducing a sparsity constraint. That work was further extended in [36], incorporating temporal information for matching visual and audio onsets.

A more recent work which focuses solely on chamber musical performances [1], explored the association of musical scores with their spatio-temporal visual locations in video recordings. First, the authors performed audio-score alignment based on chroma features and Dynamic Time Warping, therefore automatically obtaining video-score alignment. Next, they used optical flow to compute bow strokes motion velocities and correlated them with audio onsets. The video analysis consisted of fitting a Gaussian Mixture Model for player detection and computing a histogram of motion magnitudes for fine-grained localisation of a high-motion region.

Parekh et al. [20] looked for sparse motion patterns which are similar to audio activation matrices obtained with Non-negative Matrix Factorization (NMF). In particular, from the visual modality, the authors computed frame-wise average magnitude velocities of clustered motion trajectories. Then, a linear transformation which transforms the motion velocity matrix into the spectral activation matrix was used to constrain the non-negative least square cost function together with a sparsity constraint. Both NMF and the audio-motion transformation were jointly optimised. The results showed a noticeable drop in signal-to-distortion ratio (SDR) when going from duos to quartets (from 7.14dB to 0.67dB for the best method while using soft masks for reconstruction). As in [1], the performance of the method proposed in [20] decreases when separating sounds of the same instrument while addressing this problem for the first

time. Interestingly, the authors only focused on the motion component of videos ignoring other visual characteristics such as shape, color, and texture.

2.2.2 Audio-visual deep learning methods

With the outbreak of deep learning techniques, the field of audio-visual learning has received a significant boost, especially the problems formulated in unsupervised and self-supervised manners. Along this line of research there have been some works focused on representation learning with further applications in audio classification, action recognition and source localisation [18], [37], [38], [39], [40], [41], [42], [43]. Most of them combined features from two-stream networks (one sub-network for the audio and another one for the visual modality) either by concatenating them or by having an additional attention module. Some of them employed time synchrony for the samples of the same video [18], [44], while others learnt to extract features by identifying if the audio sample corresponded to a given visual data [18], [37], [40]. More recent work also focused on the usage of audio for distilling redundant visual information to reduce computational costs [41].

Different objective functions such as cross-entropy [38], [39], Kullback–Leibler divergence [37], [41], contrastive [18] or triplet [40] losses were exploited in audio-visual deep learning. Distinctively, Korbar et al. [18] used curriculum learning by first training the network with easy examples (correspondence was defined as being sampled from the same video) and then with hard/superhard examples (correspondence was defined as time-synchrony with/without time shift within the same video).

At the same time, the field of visually assisted source separation has emerged [10], [44], [45], [46], [47], in particular, with explicit focus on musical data [2], [7], [8], [9], [10]. Starting with capturing only visual appearance features [7], [8], [9], [10] there is a shift towards capturing and integrating motion data [2].

To combine the data obtained from different modalities, commonly used approaches include late fusion [47], conditioning at the bottleneck via tile-and-multiply [9], concatenation [45], attention mechanism [2], [7], and Feature-wise Linear Modulation (FiLM) conditioning [2], [48] (more details in Section 2.3). In the present work, we also analyse different ways to combine audio and visual information and extend prior work for multiple and variable number of sources in the mixture.

It's also worth noting that two antecedent works in audio-visual source separation explored approaches that separate one source at a time, which can be applied for estimating multiple and variable number of sources [9], [10]. However, they have only been trained on *artificial* mixtures of up to 4 sources and *real* mixtures of 2 sources. The separation enhancement scheme proposed in [10] consisted of extracting one source at a time from a residual audio mixture while considering maximum visual energy at every step, which follows the idea proposed in [29]. Authors trained the network with mixtures of 2 and 3 instruments, and tested it on mixtures of up to 5 instruments.

Concurrently, the idea of co-separation has been proposed in [9]. The method consisted of guiding source separation by integrating visual features of a detected musical

instrument at the bottleneck of the primary U-Net, while the training was done using mix-and-separate approach with a combination of separation and consistency losses. The latter is defined as a cross-entropy loss between ground truth instrument labels and the predictions obtained with an additional classifier on the preliminary separated sources.

2.3 Conditioned source separation

In the previous section we reviewed existing research in source separation combining information from visual and audio modalities. This approach can be reformulated as audio source separation *conditioned* on visual information. We observe that, while there are several strategies of data fusion (i.e. concatenation or co-processing), another possibility is to modulate activations of a primary audio network by a context vector extracted from another modality. This technique is known as **Feature-wise Linear Modulation (FiLM)** [48]. The conceptual idea of FiLM conditioning is simple: it takes a set of learned features and scale and shift them accordingly to a context vector. Scaling and shifting parameters (γ, β) are learned based on an input context vector $\mathbf{c}$ by an arbitrary function f which is called FiLM-generator:

$$(\gamma, \beta) = f(\mathbf{c}). \quad (4)$$

The learned parameters modulate a neural network's activations F_i via a feature-wise affine transformation with i the feature map index:

$$FiLM(F_i | \gamma_i, \beta_i) = \gamma_i F_i + \beta_i. \quad (5)$$

Other studies considered *weak conditioning* in source separation using only labels of target sources [19], [49] in contrast to *strong conditioning* where the context vector could be available frame-wise [50], [51]. The employed weak label conditioning techniques include FiLM [19] and tile-and-multiply [49]. For strong conditioning, a binary vocal activity vector and vocals magnitude vector have been used for singing voice separation with attention mechanism [51].

Later, the idea has been explored in the context of universal source separation with conditioning on classification embeddings [50]. First, the method extracts the context embeddings with the classification network, then upsample and normalise them, which is followed by conditioning of the primary source separation network either by concatenation with the network's activations or gating the activations by the embeddings. Another work goes along this line and trained a source separation model based solely on weak labels [52]. The method consists of training a classifier network and using the classification loss (with an additional constraint for the estimated sources to sum to the mixture) as the objective function for separation.

We find various strategies to integrate side information, and different modules of the network being conditioned. However, most of the studies injected the context vector at the bottleneck of encoder-decoder architecture with a rare exception of early fusion in [50]. The same authors [50] reported that integration of the context vector at every layer of the primary network leads to overfitting.

3 APPROACH

In this work, we study the effect of integrating two types of context information at different locations of the network, while keeping the architecture fixed and simple. As context information we explore: i) labels, that indicate the presence or absence of each type of instrument in the mixture, and ii) visual features.

To the best of our knowledge, there is a lack of public datasets in the format of audio-visual music recordings with ground-truth individual sources. Thus, we use a mix-and-separate approach for training, such that every mixture is generated on the fly and, therefore, unique. To create a mixture, we take the following steps: (1) we sample an arbitrary subset of instruments; (2) we subsequently pick a random segment from one of the audios of each instrument category; and (3) we scale values in each segment to the $[-1, 1]$ range in order to adjust loudness, then sum the segments and normalize the sum by the number of instruments to avoid clipping. Given a magnitude spectrogram of the mixture, our network learns to predict K real-valued masks $\hat{M}_i$, one mask per potential instrument present in the mixture (we use $K = 13$ different instruments in our experiments, see Figure 1a U-Net and Figure 2). Each output mask is associated to a certain type of instrument, and their order is fixed.

3.1 U-Net and Multi-Head U-Net baselines

We adopt two simple U-Net models as the baseline architectures, given that U-Net has been extensively used and demonstrated good performance [2], [7], [9], [26], [27]. As the focus of this work is on studying the effect of different types of conditioning, we leave for future research the analysis of different source separation networks.

U-Net [53] is an encoder-decoder architecture with *skip connections* such that activations of every i -th layer of the encoder are concatenated with activations of $(N - i)$ -th layer of the decoder, which can be considered as a light form of conditioning by itself. Following [2], [7], we have chosen one of the architectures they propose and set the number of layers to $N = 6$. We employ two variants of the architecture, namely: (a) a baseline U-Net architecture as pictured in Figure 1a which outputs 13 masks after the last upconvolutional layer, and (b) Multi-Head U-Net (MHU-Net) [27] as pictured in Figure 1b which has a single shared encoder and 13 decoders, where each dedicated decoder yields a mask for its corresponding instrument.

Audio is resampled at 11025 Hz before preprocessing in order to speed up computation, based on [7], where the authors argue that most perceptually important frequencies of instruments are preserved and audio quality is only slightly reduced. We use Hann window, and the STFT is computed for every segment of approx. 6 seconds (65535 audio samples) with window size of 1022 (about 93ms) and hop size of 256, which results in a matrix of 512×256 STFT bins. All those parameters are taken from [7] and [9] (the same hop length and approx. the same window length in time were also used in [10]). Next, we study a few preprocessing strategies over the STFT representation, including linear and log-sampled frequency scale for the

Fig. 1. Summary of architectures and conditioning types. There are two variants for the source separation architecture: (a) U-Net which outputs 13 masks at the last upconvolutional layer, (b) Multi-Head U-Net with one shared encoder and 13 specialized decoders which output one mask each. (c) There are several choices for U-Net conditioning: three types of FiLM conditioning and multiplicative conditioning of the output masks.

STFT, as well as log-scale and dB-scale with normalisation for the STFT magnitude values as discussed in Section 4.4.

The choice of loss functions is dependent on the type of mask. For binary masks, we compute binary cross entropy (BCE) loss:

$$\mathcal{L}^b = - \sum_{i=1}^K \sum_{(\tau, \omega)} (\lambda M_i^{ib}(\tau, \omega) \log(\hat{M}_i(\tau, \omega)) + (1 - M_i^{ib}(\tau, \omega)) \log(1 - \hat{M}_i(\tau, \omega))), \quad (6)$$

where $M_i^{ib}(\tau, \omega)$ and $\hat{M}_i(\tau, \omega)$ represent, respectively, ground truth and predicted binary mask values at frequency bin ω and time frame τ , K is the maximum number of sources, and λ is a positive (fixed) weight which is used to compensate for the class imbalance in the mask values.

For ratio masks we employ L_2 loss:

$$\mathcal{L}^r = \sum_{i=1}^K \|M_i^{ir} - \hat{M}_i\|_2^2, \quad (7)$$

where M_i^{ir} and $\hat{M}_i$ denote, respectively, ground truth and predicted ratio masks.

3.2 Conditioned U-Net

In this section we describe the conditioning strategies and the types of context data (Figure 2b) which we use in our Conditioned U-Net architecture (Figure 1c).

3.2.1 Weak label conditioning

We study *weak conditioning* for source separation which means that instrument labels are available at the level of individual recordings. They indicate the presence or absence of each instrument in the mix, which is encoded in a binary indicator vector $\mathbf{c}_l \in \{0, 1\}^K$ where K is the total number of instrument classes considered. In all network variants, the output is a set of K masks (Figure 2a), and we denote by $\hat{M}_i$ the mask associated to the i -th instrument type.

We use $\mathbf{c}_l$ as a conditioning context vector and compare three types of FiLM conditioning: introduced (1) at the

(a) The main separation network receives as input the spectrogram of the mixture and a context vector. It outputs K masks.

(b) Types of context vectors: binary (left), visual (center), and visual-motion (right). The inputs in each case are: binary vector (left), one frame per instrument (center), a set of frames per instrument (right).

Fig. 2. Diagrams showing the inputs, outputs and types of conditioning vectors.

bottleneck, (2) at all encoder layers, and (3) at the final decoder layer (see Figure 1c). More formally, for each layer j (no matter at which location the FiLM is used) we have activations or embeddings: $\mathbf{a}^{(j)}$, and the conditioning is as follows:

$$(\gamma_j, \beta_j) = f_j(\mathbf{c}_l) \quad (8)$$

$$\hat{\mathbf{a}}^{(j)} = \gamma_j \mathbf{a}^{(j)} + \beta_j \quad (9)$$

where f_j is a single fully connected layer.

Furthermore, we explore simple multiplicative conditioning with the binary indicator vector:

$$\hat{M}_i = \mathbf{c}_l[i] \tilde{M}_i, \quad (10)$$

where $\mathbf{c}_l[k]$ is the i -th component of the context vector and $\tilde{M}_i$ is the i -th preliminary mask as predicted by (MH)U-Net.

3.2.2 Visual conditioning

In the case of visually-informed source separation, we consider both static characteristics and motion-aware conditioning. Nonetheless, we would like to note that learning temporal information from videos is a challenging task which is still under research. Therefore, visually-informed methods mostly use a single frame for conditioning [7], [8], [9], with some exception of dense trajectories [1], and deep-learned dense trajectories [2].

Alike [9], [20] we assume that rough spatial location of each source is given (e.g. it can be obtained by a segmentation or human detection algorithm). Keeping this assumption in mind, we use uncropped frames from individual videos for training and evaluation (since videos in the Solos dataset contain a single instrument). In a real life scenario (e.g. for testing) we use a bounding box around every player. In the experiments with the URMP dataset we manually annotated the bounding boxes.

For visual context conditioning, we take a single video frame corresponding to the beginning of the audio source sample. We use a pretrained ResNet-50 [54] to extract a visual feature vector of size 2048 for every present source, and then concatenate them, obtaining a visual context vector $\mathbf{c}_v$ of size $K' \times 2048$ where K' is the maximum number of sources in the mixture ($K' \leq K$). The context vector for the unavailable sources is set to all zeros. To obtain a final conditioning vector, we first pass the visual feature vectors

through an adaptive max pooling and obtain a conditioning of size $K' \times 1024$. Then, we take a per-channel maximum over all K' frames and obtain a single 1024-D conditioning vector. As for the case of weak label conditioning, we compare three alternatives for the FiLM conditioning (see Figure 1c). Additionally, we conduct an experiment which is similar to label multiplicative conditioning, but instead of using labels we use a vector of instrument probabilities. Class probabilities are obtained by adding an additional fully-connected layer with a softmax activation after the context vectors. This experiment is denoted later in Section 5.3 as Final-multiply.

For visual-motion conditioning, we first extract visual feature vectors with the pretrained ResNet-50 at a fixed frame-rate within the video segment corresponding to the selected 6s segment of audio. We then pass the obtained sequence of vectors through a small uni-directional LSTM network with 1024 hidden units, as in [41], with the aim of capturing motion characteristics while keeping visual information. We take the last LSTM hidden state resulting in a motion context vector $\mathbf{c}_m$ of size 1024. Due to the large computational cost, and based on the results of the ablation study (Section 5.3), we only report this approach with FiLM conditioning at the bottleneck and the final layer of audio U-Net. Alternatively, we analyse another strategy which consists of directly max-pooling visual features extracted from all frames into a single conditioning vector.

4 EXPERIMENTS

In what follows, we thoroughly evaluate the proposed method on various setups. In particular, we compare the different conditioned networks with respect to several performance metrics.

4.1 Dataset

In our experiments we use two multimodal datasets of musical performances: the recently introduced Solos dataset [55] for training and evaluation, and the URMP dataset [11] for testing (Table 1 provides a summary of both datasets).

The URMP dataset² consists of 44 arrangements (of which 11 are duets, 12 are trios, 14 are quartets, and 7 are

2. <http://www2.ece.rochester.edu/projects/air/projects/URMP.html>

quintets) with a total duration of 1h 14min. Each source track was recorded separately with an external coordination, and the final mixes were assembled afterwards. The instrumentation is a typical one for chamber and orchestral music, and includes families of instruments such as strings (violin, viola, cello and double bass), woodwinds (flute, oboe, clarinet, bassoon, saxophone), and brass (trumpet, horn, trombone, tuba). The dataset is constructed to reflect the complexity of the musical world where the same instrument within a section can appear more than once. For each piece, the dataset provides the musical score in MIDI format, the high-quality individual instrument audio recordings and the videos of the assembled pieces. For that, all individual visual contents were composited in a single video with a common background where all players were arranged at the same level from left to right.

As in this work we only tackle the problem of separating sources of different instruments, we mix source tracks of the same instrument within the same piece and consider the resulting mix as a single source. For example, for a string quartet (which consists of 2 violins, viola and cello), we mix the two source tracks of the violins which results in a corresponding “trio” where two violins are considered as a single source. Also, we remove four pieces (02_Sonata_vn_vn, 04_Allegro_fl_fl, 05_Entertainer_tpt_tpt, 06_Entertainer_sax_sax) from the dataset as they are duets of the same instrument. After this preprocessing, we are left with 12 duets, 20 trios and 8 quartets in the final set.

The Solos dataset³ consists of 755 YouTube videos of solo musical performances of the same 13 instruments categories as of the URMP dataset. It has a total duration of about 66 hours. A major part of the dataset are audition performances which ensures, together with manual and semi-automatic checking, a good quality of audio and video. We use this data with a mix-and-separate strategy and complement it with an evaluation of the trained models in the URMP dataset, which allows proper evaluation on real-world mixtures. Although most of the previous works on audio-visual learning-based source separation have been trained on the MUSIC dataset introduced in [7], it lacks 6 out of the 13 instrument classes present in the URMP dataset (namely, viola, double bass, oboe, bassoon, horn, trombone). That is why we choose to train on Solos.

4.2 Metrics

Several studies indicate that widely-adopted source separation metrics such as signal to distortion ratio (SDR), signal to inference ratio (SIR), and signal to artifacts ratio (SAR) [56] do not always agree with human perception [7], [9], [35], [57]. Moreover, as brought out in [35], an increment of noise or interferences in the separated source produces an increment of the SAR value. Recently, *scale-invariant* and *scale-dependent* SDR (SI-SDR, SD-SDR) have been proposed [35] as better alternative metrics. For completeness, we report all these measures.

Unlike previous works [2], [7], [9], [10], our method produces *sparse* outputs since many predicted sources are expected to be silent. However, all the above metrics are ill-defined for silent sources and targets. To address this issue,

we also compute cumulative *predicted energy at silence* (PES) as proposed in [58] and later refined in [51]. For SI-SDR and SD-SDR larger values indicate better performance, for PES smaller values indicate better performance. For numerical stability of log function, in our implementation we add a small constant $\epsilon = 10^{-9}$ which results in a lower boundary of the metrics to be -80 dB.

4.3 Training and implementation details

Our U-Net is composed of six blocks in the encoder and six blocks in the decoder. Each encoder block consists of a convolutional layer followed by batch normalisation with an optional conditioning layer (for FiLM-encoder conditioning), and LeakyReLU nonlinearity. We use convolutions of 4×4 with stride two and padding one. The number of filters in the initial layer is set to 16, and then doubled for each consecutive layer until it reaches 512 in the 6th block. A decoder block consists of a bilinear upsampling layer with a scale factor of two, a convolutional layer, batch normalisation, ReLU nonlinearity, and a dropout layer with $p = 0.2$. The convolutional filters in the decoder blocks are of size 3×3 with stride one and padding one.

The network is trained for 1000 epochs with a batch size of 32, Adam optimizer, an initial learning rate of 10^{-5} which is halved after 25k iterations with no improvement on the validation set.

We have also analysed the use of a curriculum learning strategy. It consists of starting the training with only mixtures of 2 sources, and gradually increasing the maximum number of sources up to 7. The increment is carried out if the validation loss does not decrease for 10k iterations. Some examples where curriculum learning has proven to be beneficial for training audio-visual networks are [2] for source separation, and [18] for audio-visual temporal synchronisation.

For training and evaluation we utilize the mix-and-separate procedure by creating artificial mixtures from individual videos of Solos. Every training sample has an arbitrary number of sources in the mixture, with an upper bound of the maximum number of sources to appear simultaneously is set to 7. For testing, we use real mixtures from the URMP dataset which have at most 4 simultaneously sounding sources.

The implementation takes advantage of various Python libraries and uses PyTorch as a deep learning framework. The code is made publicly available⁴.

4.4 Baseline ablation study

In preparation for conditioned source separation analysis and to define the optimal hyperparameters of the baseline U-Net architecture as described in Section 3.1, we conduct a series of ablation experiments. We examine the following set of hyperparameters: (1) linear vs. log frequency scale for the STFT representation, (2) binary vs. ratio masks estimation, (3) data augmentation with normally-distributed noise, (4) log vs. dB-normalised scale for the STFT values, (5) the use of curriculum learning, and (6) the effectiveness of Multi-Head U-Net vs. vanilla U-Net.

4. <https://github.com/Veleslavia/conditioned-u-net/tree/master/code>

3. <https://www.juanmontesinos.com/Solos/>

Category	URMP				Solos	
	# Recordings	# Pieces	(# duets, # trios, # quartets)	Mean duration	# Recordings	Mean duration
Violin	11	20	(5, 12, 3)	1:51	66	6:16
Viola	7	12	(1, 8, 3)	1:52	55	5:31
Cello	10	11	(3, 7, 1)	1:52	134	7:21
DoubleBass	3	4	(0, 1, 2)	2:28	58	8:53
Flute	8	11	(3, 4, 4)	1:50	48	4:00
Oboe	3	6	(1, 2, 3)	1:58	53	5:45
Clarinet	8	10	(2, 5, 3)	1:44	49	3:23
Bassoon	2	3	(0, 0, 3)	1:25	56	5:08
Saxophone	8	9	(2, 5, 2)	1:31	45	2:42
Trumpet	11	12	(5, 5, 2)	2:07	50	1:14
Horn	5	5	(0, 3, 2)	2:02	50	5:11
Trombone	8	9	(2, 5, 2)	2:01	50	5:03
Tuba	3	5	(0, 3, 2)	1:47	41	2:49
TOTAL	87	40	(12, 20, 8)	1:53	755	5:16

TABLE 1

Statistics of the datasets used: URMP and Solos. Notice that, as commented in Section 4.1, we consider different instances of the same instrument as the same source, that is why the resulting number of pieces, duos, trios, and quartets varies from the original dataset (e.g. there are no quintets left since in all of them at least one of the instruments is repeated).

Method	Exp. ID	W	SI-SDR $\uparrow$	SD-SDR $\uparrow$	PES $\downarrow$	SDR $\uparrow$	SIR $\uparrow$	SAR $\uparrow$
IRM	U	0	13.1 ± 5.4	12.7 ± 6.4	n/a	11.8 ± 4.3	19.9 ± 5.6	13 ± 4.3
input mix	L	0	-3.7 ± 5.7	-3.7 ± 5.7	18.2 ± 4.2	-3.5 ± 4.8	-3.2 ± 4.9	18.1 ± 11.2
No filtering	1	0	-2.1 ± 6.2	-15.2 ± 10.0	-8.1 ± 5.9	-0.8 ± 5.8	0.6 ± 6.8	10.9 ± 3.2
Wiener (1 it.)	1	1	-3.9 ± 7.5	-16.0 ± 16.1	-15.1 ± 10.1	-1 ± 6.2	1.5 ± 7.7	7.7 ± 3.8
Wiener (2 it.)	1	2	-7.6 ± 10.7	-21.4 ± 24.1	-25.4 ± 16.8	-2.6 ± 7.1	1.7 ± 8.5	5 ± 5.6
Exp. 1 + CL	2	2	-8.2 ± 10.9	-24.1 ± 25.1	-26.2 ± 17.8	-3.1 ± 7.1	1.4 ± 8.5	4.6 ± 5.7
Exp. 1 + BM	3	2	-9.7 ± 15.2	-21.6 ± 27.4	-36.2 ± 18.4	-3.8 ± 8	0.1 ± 7.1	4.9 ± 7.5
Exp. 1 + noise	4	2	-7.6 ± 9.2	-21.7 ± 20.4	-21.7 ± 15.6	-3.1 ± 6.9	0.2 ± 8.3	6.0 ± 5.3
Exp. 1 + log-normalise	5	2	-10.1 ± 11.6	-25.3 ± 24.9	-27.0 ± 17.0	-3.9 ± 7.6	0.7 ± 9.5	4.6 ± 6.1
Exp. 1 + linear STFT	6	2	-7.4 ± 9.9	-17.5 ± 19.8	-24.4 ± 15.8	-2.3 ± 6.9	1.5 ± 8.4	5.7 ± 5.2
Exp. 1 + MHU-Net	7	2	-8.1 ± 9.1	-23.7 ± 20.4	-20.9 ± 15.8	-3.3 ± 6.4	0.1 ± 8.2	5.8 ± 4.5

TABLE 2

Ablation studies results for the URMP dataset with no conditioning. The first two rows indicate two references: ideal ratio masks (IRM, U states for the upper bound), and the usage of the input mixture as a predicted source (input mix, L states for the lower bound). Experiment IDs are explained in Table 5. The column W indicates the number of iterations of Wiener filtering used in post-processing. The best results among all variants with 2 iterations of Wiener filtering are highlighted in **bold**. Results that are **not statistically significant** ($p < 0.05$) w.r.t. Exp. 1 with 2 iterations of Wiener filtering are dotted. CL denotes the usage of curriculum learning and BM states for binary masks.

5 RESULTS AND DISCUSSION

5.1 Ablation studies

We report the metrics obtained by unconditioned U-Net models in the ablation study in Table 2, and the full list of hyperparameters is given in Appendix A. The experiments can be matched by the experiment ID. We also provide two reference metrics, the upper bound separation quality (U) with ideal ratio masks (IRM), and the mixture metrics (L) which reproduce the input mixture at every possible output source. Results that are **not** statistically significant w.r.t. Exp. 1 ($p < 0.05$) are dotted.

In practice, we noticed that, while training with ratio masks, the to-be-silent output sources eventually happen to be an original mixture with a lowered volume. Therefore, as it is common practice in audio source separation, we have applied a post-processing based on some iterations of the Wiener filter [59]. We indicate the number of iterations of Wiener filtering applied in the third column (W) of Table 2. The effect of the post-processing iterations can be observed: although some of the metrics are degraded, the improvement in PES and SIR is considerable. The metric degradation can be seen mainly in Table 2, in the second horizontal block. Figure 5 allows deeper analysis of quality changes, which is generic for post-processing in all experiments. In spite of the fact that the post-processed results sound more filtered and have artefacts, we can perceive a

notable separation improvement when it comes to interferences. We believe that for multi-source separation with many silent source, the practical trade-off between quality degradation and separation improvement offered by post-processing is worth considering. In a real case scenario, the final choice would probably depend on the application. For example, for remixing it would be beneficial to preserve the original quality, but fewer interferences could be more desirable if the source separation is used as a preliminary step for automatic transcription. In this study we try to find a balance between different performance metrics and different perceptual aspects of separated audio. Therefore, in all following experiments we apply two iterations of Wiener post-processing.

As mentioned in Section 4.3, we have analysed the use of a curriculum learning (CL) strategy (Exp. 2) but as it can be observed it is only slightly beneficial in terms of PES while comparing with Exp. 1. We have also used binary masks instead of ratio masks (Exp. 3), the results are worse in general but the silent sources are better recovered which is resulted in lower PES (-10.8dB comparing to Exp. 1). Another aspect we have observed is that augmenting the input data with normally-distributed noise (Exp. 4) does not improve separation performance (except for SAR). However, we are of the view that the problem can benefit from other augmentation techniques, such as the ones described

in [31], [60]. This is left for future research.

Results from Exp. 5 show that using normalised STFT is beneficial and Exp. 6 reveals that working with a linear-scale STFT gives overall better results. All measures in Exp. 6 have improved with respect to a log-scale STFT (Exp. 1) with the exception of PES and SIR which are barely degraded. In particular, the improvements in SD-SDR and SAR are statistically significant. Finally, the use of a multi-decoder architecture (Exp. 7), which requires higher computational cost, does not improve the metrics (except for SAR).

In summary, the optimal parameters we found are dB-normalised and linear-scale STFT as input, a single decoder that predicts ratio masks, and Wiener post-processing. We have opted out of augmenting the input with normally-distributed noise and using curriculum learning. This configuration will be used in the following sections as baseline for the conditioning experiments.

5.2 Conditioning on labels

We further study weak label conditioning of the single-decoder U-Net model. We provide results for four conditioning schemes as described in Section 3.2.1. The summary of weak label conditioned source separation is shown in Table 3.

We observe that the best performance in terms of SI-SDR, SDR, and SAR is obtained with FiLM conditioning in the encoder (Exp. 9, Table 3), but it also leads to high PES, even worse than in the case of no post-processing in the ablation study (Exp. 1, $W = 0$, Table 2). On the other hand, the best results in terms of SD-SDR and SIR are achieved with the multiplicative conditioning on the output masks (Exp. 11, Table 3).

Figure 3 shows the performance measured by SI-SDR, SD-SDR and PES in Exp. 11 for each instrument in the URMP dataset. The results emphasize the fluctuations between the instruments. As it can be observed, the worst SI-SDR and SD-SDR values correspond to clarinet, while the best three instruments in terms of SI-SDR are cello, trumpet, flute and double bass. Tuba is also the best instrument with respect to PES.

On the other hand, Figure 4 presents the SDR, SIR, and SAR metrics with respect to the number of sources in the mixture for two different models: the baseline network without conditioning (Exp. 6) and the same network with Label-multiply conditioning (Exp. 11). The improvements are consistent among the metrics and the number of voices in the mixture.

5.3 Conditioning on visual information

We compare visually conditioned U-Net with its corresponding non-conditioned and label conditioned baselines.

Table 4 shows the performance of single-frame visually conditioned U-Net given the same FiLM locations as in the label conditioning case. It also indicates the results of conditioning by a visual-motion context vector learned from 5 frames per segment (with the frame rate set to 5/6 fps). Lastly, we report the results for the Sound-of-Pixels (SoP) [7] method. SoP-unet7 states for the original method trained on the Music dataset published in [7]. We used the officially provided weights and evaluated the model on the URMP

Fig. 3. Exp. 11 per-instrument boxplots for the URMP dataset. Note that the x axis scale limits vary from metric to metric. The principal reason of the difference between SI-SDR and SD-SDR is that SD-SDR accounts for the volume changes.

Fig. 4. Exp. 6 (no conditioning) vs. Exp. 11 (weak label-conditioning, multiply) SDR, SIR and SAR boxplots for the URMP dataset with respect to the number of sources in the mixture.

dataset. SoP-unet7-ft indicates the version which was fine-tuned on the Solos dataset. SoP-unet5-Solos accounts for a model with 5 blocks in U-Net which is trained from scratch. In all SoP networks, both visual and audio networks are trained simultaneously while in our conditioning experiments visual network is fine-tuned in all experiments with a single-frame conditioning and is frozen in all experiments with a multi-frame conditioning.

The metrics show that the best results both in terms of SI-SDR and SDR are attained by multiplicative conditioning

Fig. 5. Effect of post-processing with 1 and 2 iterations of Wiener filtering for a test sample from Exp. 11 results. This is an example of separated cello with interferences from violin and viola voices. We can observe that, while audio quality is degrading and the separated part becomes more filtered, the interferences are also greatly reduced. We use linear STFT scale and visualise frequencies up to 5.5 kHz.

Method	Exp. ID	W	SI-SDR $\uparrow$	SD-SDR $\uparrow$	PES $\downarrow$	SDR $\uparrow$	SIR $\uparrow$	SAR $\uparrow$
w/o conditioning	6	2	-7.4 $\pm$ 9.9	-17.5 $\pm$ 19.8	-24.4 $\pm$ 15.8	-2.3 $\pm$ 6.9	1.5 $\pm$ 8.4	5.7 $\pm$ 5.2
FiLM-bottleneck	8	2	-5.1 $\pm$ 6.8	-19.2 $\pm$ 10.7	-7.1 $\pm$ 9.9	-1.3 $\pm$ 6.0	1.9 $\pm$ 8.2	6.6 $\pm$ 2.9
FiLM-encoder	9	2	-4.2 $\pm$ 6.4	-19.7 $\pm$ 8.8	-6.9 $\pm$ 8.8	-1.2 $\pm$ 5.8	1.1 $\pm$ 7.5	8.5 $\pm$ 3.0
FiLM-final	10	2	-4.3 $\pm$ 6.3	-17.3 $\pm$ 9.1	-9.3 $\pm$ 11.2	-1.6 $\pm$ 5.6	0.6 $\pm$ 7.3	8.2 $\pm$ 3.0
Label-multiply	11	2	-7.2 $\pm$ 9.8	-15.3 $\pm$ 15.5	-18.6 $\pm$ 20.3	-1.9 $\pm$ 6.8	2.4 $\pm$ 8.8	5.1 $\pm$ 4.1

TABLE 3

Conditioned U-Net with Labels (URMP metrics). The best results are highlighted in **bold**. Results that are **not statistically significant** w.r.t. Exp. 6 ($p < 0.05$) are dotted.

at the final layer with a context vector extracted from a single visual frame (Exp. 16). Similar results are obtained using conditioning with visual-motion context vector (Exp. 13). The highest SAR among all visually conditioned U-Nets is achieved by a network FiLM-ed at the final layer (Exp. 18), while the latest also delivers comparable results in SI-SDR, SDR, and SIR to the best performing Final-multiply (Exp. 16). Surprisingly, in case of visual conditioning, none of the models outperforms the baseline network without conditioning in terms of SD-SDR, PES, and SIR. Finally, compared to the SoP model (including its fine-tuned version on Solos), all models with visual conditioning and the baseline without conditioning give better SI-SDR values. Most of those models outperform SoP as well in terms of SDR and SIR.

5.4 Unsuccessful attempts

We would like to report several strategies which did not improve source separation performance in our experiments.

In one of the experiments, we used L_2 loss while directly predicting spectrogram values instead of using the masking-based approach. However, the network failed to converge. We hypothesise that this behaviour accounts for the higher complexity of the spectrograms and the sparsity of the outputs. Another potential cause might be the gradients magnitudes that are not bounded as in the masking-based approach, therefore affecting the training dynamics of the network. The potential causes could be investigated in the future.

We also had an unsuccessful attempt to employ multi-task learning in order to further regularize the embedding space. In these experiments we jointly optimised classification and separation losses trying to predict which instruments are present in the mixture using the bottleneck U-Net features as an input for a small classifier consisting of a single fully-connected layer. While generally converging, the classification and separation performance were lower than the results of stand-alone models.

5.5 Discussion

In our experiments we observe that the use of additional information (apart from the audio signal) can be beneficial for the separation performance. However, integrating domain knowledge might be challenging and has to be done carefully. In this work we try to go beyond the classical “divide-and-conquer” approach in which data representations are learned separately and aggregated at the final stage of the separation pipeline. Nevertheless, our experiments show that most notable performance improvement can be achieved by multiplying learned masks by ground truth

instrument labels or probabilities taken from visual data. Merging a strong label signal into the encoder seems to be helpful but also leads to overfitting. From the results we observe that U-Net conditioned on the visual context vector performs worse in terms of SIR than the unconditioned version. A possible explanation for this observation could be due to the confusions separating musical instruments from the same family (like viola and violin which have a strongly similar visual appearance). Those confusions may result in more interferences and mispredictions when both instruments are present in the mixtures, which is a common case for the URMP dataset. Indeed, as shown in Table 1, violin and viola appear in great part of the duos, trios, and quartets of URMP dataset.

By inspecting the results obtained by the Sound-of-Pixels method, we highlight the importance of taking the source separation problem in the real-world scenario, as the method was previously tested in mix-and-separate settings and the reported results had an average SDR of 8dB. Our results demonstrate the demand for the testing on the real mixtures rather than using the mix-and-separate approach. Notably, even 5-blocks Sound-of-Pixels trained on Solos performs better than 7-blocks Sound-of-Pixels trained on Music. Fine-tuning of the original Sound-of-Pixels model allows to improve the quality of source separation for 1.2dB in SDR which also indicates the need of enlarging the datasets and enhancing their quality.

Even though from the literature we know that source separation can benefit from integrating the motion information [1], [2], [20], we would like to note that all aforementioned methods use complex pre-processing in order to extract reliable motion features, which brings attention to the problem of closing the gap between motion and audio representations. Another fact that should be noted is that all sources of information should be correctly combined, preserving synchrony between them. While for single-frame visual and weak label information it is not so important, for temporal data such as motion, pitch, and musical scores it may become a crucial aspect for successful conditioning. Consecutively, a different baseline source separation architecture, such as an RNN-based network, may improve on the current results due to its sequential nature which better preserves time-domain information.

Taking into consideration the above-mentioned observation, we can note that the U-Net architecture may be a limitation of our study, and the results may be different for other baseline architectures.

Given that the best results in terms of different metrics are achieved by using different setups (e.g. binary and ratio masks), we would like to emphasise that a further enhancement can be obtained by having the best of both

Method	Frames	Exp. ID	W	SI-SDR $\uparrow$	SD-SDR $\uparrow$	PES $\downarrow$	SDR $\uparrow$	SIR $\uparrow$	SAR $\uparrow$
w/o conditioning	0	6	2	-7.4 ± 9.9	-17.5 ± 19.8	-24.4 ± 15.8	-2.3 ± 6.9	1.5 ± 8.4	5.7 ± 5.2
FiLM-encoder	1	12	2	-12.1 ± 13.9	-34.9 ± 29.7	-21.7 ± 22.3	-5.1 ± 7.9	-0.3 ± 8.1	3.3 ± 6.9
FiLM-bottleneck	1	17	2	-5.1 ± 7.7	-17.6 ± 14.7	-16.3 ± 17.7	-2.2 ± 5.7	0.2 ± 7.0	7.2 ± 3.7
FiLM-final	1	18	2	-4.8 ± 7.3	-23.7 ± 20.2	-21.3 ± 18.7	-1.6 ± 5.6	0.8 ± 8.0	10.3 ± 5.4
Final-multiply	1	16	2	-4.1 ± 7.2	-19.0 ± 13.9	-15.4 ± 14.3	-1.5 ± 6.1	1.0 ± 8.4	9.2 ± 3.7
FiLM-final-lstm	5	13	2	-4.5 ± 6.2	-21.5 ± 11.9	-12.7 ± 13.8	-2.1 ± 5.4	-0.4 ± 6.7	9.3 ± 3.7
FiLM-bottleneck-lstm	5	14	2	-5.6 ± 7.6	-19.3 ± 15.8	-18.3 ± 16.2	-2.1 ± 5.8	0.2 ± 7.0	7.5 ± 4.1
FiLM-bottleneck-mp	5	15	2	-8.2 ± 9.0	-26.6 ± 21.4	-19.3 ± 17.7	-3.3 ± 6.3	0.5 ± 7.6	5.2 ± 5.5
SoP-unet7 [7]	3	19	0	-18.7 ± 8.9	-21.1 ± 9.4	n/a	-3.8 ± 4.0	-1.5 ± 4.7	7.6 ± 3.1
SoP-unet7-ft [7]	3	20	0	-17.5 ± 8.5	-20.3 ± 9.3	n/a	-2.6 ± 5.0	0.5 ± 6.4	6.9 ± 2.5
SoP-unet5-Solos	3	21	0	-16.9 ± 8.6	-18.7 ± 8.9	n/a	-2.9 ± 4.7	-1.7 ± 5.3	11.1 ± 6.9

TABLE 4

Results for Visually Conditioned U-Net experiments with different types of conditioning and different number of frames used (evaluated on the URMP dataset). The best results are highlighted in **bold**. Results that are **not statistically significant** w.r.t. Exp. 6 ($p < 0.05$) are dotted.

worlds. This would require an additional network that combines the estimates of binary and soft masks, as proposed in [32]. Finally, we would like to note the opportunity to surpass the current performance by employing additional constraints for the loss functions as in [9], [25], or weighting the loss values of the masks with the magnitude values of the mixture [7], [9] as it may help to avoid treating every time frequency bin equally and focus attention on the areas where most of the energy is concentrated.

6 CONCLUSION

We tackle a problem of Single Channel Source Separation for multi-instrument polyphonic music conditioned on external data. In this work we have shown that the use of extra information such as (1) binary vectors indicating the presence or absence of musical instruments in the mix and (2) visual feature vectors extracted from corresponding video frames improve the separation performance.

We also show that different types of conditioning have different effects w.r.t. the performance metrics. We have conducted a thorough study of FiLM-conditioning introduced at three possible locations of the primary source separation U-Net model. We have demonstrated that the best results in SI-SDR and SDR can be obtained with label conditioning by FiLM at the encoder or by visual conditioning at the final output mask.

The results shown in the present work indicate that the real-case scenario such as chamber quartets source separation is challenging and there is still a significant performance gap of about 13dB between the state-of-the-art separation methods and ideal ratio masks.

Potential improvements could include modifying the U-Net architecture, combining binary and soft masks to obtain a good balance between SDR and PES. Another possibility could be integrating an advanced motion analysis network and employing audio-motion synchrony for conditioning the network, and conditioning on musical scores.

APPENDIX A

HYPERPARAMETERS OF THE EXPERIMENTS

We provide the full set of model hyperparameters used in the experiments in Section 5 in Table 5. Please note that only one parameter changes between each pair of the experiments compared in Table 2. For the experiments in Section 5.3 the model parameters are set as described in Section 5.1.

APPENDIX B

PER-EXPERIMENT BAR PLOTS WITH SOURCE SEPARATION PERFORMANCE RESULTS

Figure 6 shows source separation metrics (SI-SDR, SD-SDR, PES) pictured as bar plots with mean and standard deviation for each experiment conducted in Section 5. The experiment are referenced by id as in Tables 2, 3 and 4.

ACKNOWLEDGMENT

This work was funded in part by ERC Innovation Programme (grant 770376, TROMPA); Spanish Ministry of Economy and Competitiveness under the María de Maeztu Units of Excellence Program (MDM-2015-0502) and the Social European Funds; the MICINN/FEDER UE project with reference PGC2018-098625-B-I00; and the H2020-MSCA-RISE-2017 project with reference 777826 NoMADS. We gratefully acknowledge NVIDIA for the donation of GPUs used for the experiments.

Exp. ID	STFT F-scale	STFT V-scale	model	noise	mask	loss	CL	conditioning
1	log	dB-norm	U-Net	No	Ratio	L_2	No	None
2	log	dB-norm	U-Net	No	Ratio	L_2	Yes	None
3	log	dB-norm	U-Net	No	Binary	BCE	No	None
4	log	dB-norm	U-Net	Yes	Ratio	L_2	No	None
5	log	log	U-Net	No	Ratio	L_2	No	None
6	linear	dB-norm	U-Net	No	Ratio	L_2	No	None
7	log	dB-norm	MHU-Net	No	Ratio	L_2	No	None
8	linear	dB-norm	U-Net	No	Ratio	L_2	No	FiLM-bottleneck (binary)
9	linear	dB-norm	U-Net	No	Ratio	L_2	No	FiLM-encoder (binary)
10	linear	dB-norm	U-Net	No	Ratio	L_2	No	FiLM-final (binary)
11	linear	dB-norm	U-Net	No	Ratio	L_2	No	Label-multiply (binary)
12	linear	dB-norm	U-Net	No	Ratio	L_2	No	FiLM-encoder (visual)
13	linear	dB-norm	U-Net	No	Ratio	L_2	No	FiLM-final-lstm (visual-motion)
14	linear	dB-norm	U-Net	No	Ratio	L_2	No	FiLM-bottleneck-lstm (visual-motion)
15	linear	dB-norm	U-Net	No	Ratio	L_2	No	FiLM-bottleneck-maxpool (visual-motion)
16	linear	dB-norm	U-Net	No	Ratio	L_2	No	Final-multiply (visual)
17	linear	dB-norm	U-Net	No	Ratio	L_2	No	FiLM-bottleneck (visual)
18	linear	dB-norm	U-Net	No	Ratio	L_2	No	FiLM-final (visual)
19	log	log	SoP-7	No	Binary	BCE	No	SoP-LinearCombination (visual)
20	log	log	SoP-7-ft	No	Binary	BCE	No	SoP-LinearCombination (visual)
21	log	log	SoP-5	No	Binary	BCE	No	SoP-LinearCombination (visual)

TABLE 5

Hyperparameters for experiments reported in this study and corresponding experiment IDs for Conditioned U-Net. CL denotes the usage of curriculum learning.

Fig. 6. SI-SDR, SD-SDR and PES boxplots for the experiments from Section 5. Experiments are referenced by ID.

REFERENCES

- [1] B. Li, K. Dinesh, Z. Duan, and G. Sharma, "See and listen: Score-informed association of sound tracks to players in chamber music performance videos," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 2906–2910.
- [2] H. Zhao, C. Gan, W.-C. Ma, and A. Torralba, "The sound of motions," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 1735–1744.
- [3] M. Adeli, J. Rouat, and S. Molotchnikoff, "Audiovisual correspondence between musical timbre and visual shapes," *Frontiers in human neuroscience*, vol. 8, p. 352, 2014.
- [4] M. Miron, J. J. Carabias-Orti, J. J. Bosch, E. Gómez, and J. Janer, "Score-informed source separation for multichannel orchestral recordings," *Journal of Electrical and Computer Engineering*, vol. 2016, 2016.
- [5] J. J. Carabias-Orti, M. Cobos, P. Vera-Candeas, and F. J. Rodríguez-Serrano, "Nonnegative signal factorization with learnt instrument models for sound source separation in close-microphone recordings," *EURASIP Journal on Advances in Signal Processing*, vol. 2013, no. 1, p. 184, 2013.
- [6] K. Lee, "An analysis and comparison of the clarinet and viola versions of the two sonatas for clarinet (or viola) and piano Op. 120 by Johannes Brahms," Ph.D. dissertation, University of Cincinnati, 2004.
- [7] H. Zhao, C. Gan, A. Rouditchenko, C. Vondrick, J. McDermott, and A. Torralba, "The sound of pixels," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 570–586.
- [8] R. Gao, R. Feris, and K. Grauman, "Learning to separate object sounds by watching unlabeled video," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 35–53.
- [9] R. Gao and K. Grauman, "Co-separating sounds of visual objects," in *Proceedings of the IEEE International Conference on Computer*

- Vision*, 2019, pp. 3879–3888.
- [10] X. Xu, B. Dai, and D. Lin, “Recursive visual sound separation using minus-plus net,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 882–891.
 - [11] B. Li, X. Liu, K. Dinesh, Z. Duan, and G. Sharma, “Creating a multitrack classical music performance dataset for multimodal music analysis: Challenges, insights, and applications,” *IEEE Transactions on Multimedia*, vol. 21, no. 2, pp. 522–535, 2019.
 - [12] E. Kidron, Y. Y. Schechner, and M. Elad, “Pixels that sound,” in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, vol. 1. IEEE, 2005, pp. 88–95.
 - [13] J. R. Hershey and J. R. Movellan, “Audio vision: Using audio-visual synchrony to locate sounds,” in *Advances in neural information processing systems*, 2000, pp. 813–819.
 - [14] W. Guo, J. Wang, and S. Wang, “Deep multimodal representation learning: A survey,” *IEEE Access*, vol. 7, pp. 63 373–63 394, 2019.
 - [15] T. Baltrušaitis, C. Ahuja, and L.-P. Morency, “Multimodal machine learning: A survey and taxonomy,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 2, pp. 423–443, 2018.
 - [16] D. Ramachandram and G. W. Taylor, “Deep multimodal learning: A survey on recent advances and trends,” *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 96–108, 2017.
 - [17] D. Michelsanti, Z.-H. Tan, S.-X. Zhang, Y. Xu, M. Yu, D. Yu, and J. Jensen, “An overview of deep-learning-based audio-visual speech enhancement and separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1368–1396, 2021.
 - [18] B. Korbar, D. Tran, and L. Torresani, “Cooperative learning of audio and video models from self-supervised synchronization,” in *Advances in Neural Information Processing Systems*, 2018, pp. 7763–7774.
 - [19] G. Meseguer-Brocal and G. Peeters, “Conditioned-u-net: Introducing a control mechanism in the u-net for multiple source separations,” in *20th International Society for Music Information Retrieval Conference (ISMIR)*, 2019.
 - [20] S. Parekh, S. Essid, A. Ozerov, N. Q. Duong, P. Pérez, and G. Richard, “Guiding audio source separation by video object information,” in *Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2017 IEEE Workshop on. IEEE, 2017, pp. 61–65.
 - [21] J. J. Carabias-Orti, T. Virtanen, P. Vera-Candeas, N. Ruiz-Reyes, and F. J. Canadas-Quesada, “Musical instrument sound multi-excitation model for non-negative spectrogram factorization,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 6, pp. 1144–1158, 2011.
 - [22] A. Ozerov and C. Févotte, “Multichannel nonnegative matrix factorization in convolutive mixtures for audio source separation,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 3, pp. 550–563, 2009.
 - [23] T. Virtanen, “Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria,” *IEEE transactions on audio, speech, and language processing*, vol. 15, no. 3, pp. 1066–1074, 2007.
 - [24] P. Chandna, M. Miron, J. Janer, and E. Gómez, “Monaural audio source separation using deep convolutional neural networks,” in *International Conference on Latent Variable Analysis and Signal Separation*. Springer, 2017, pp. 258–266.
 - [25] S. Wisdom, J. R. Hershey, K. Wilson, J. Thorpe, M. Chinen, B. Patton, and R. A. Saurous, “Differentiable consistency constraints for improved deep speech enhancement,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 900–904.
 - [26] A. Jansson, E. Humphrey, N. Montecchio, R. Bittner, A. Kumar, and T. Weyde, “Singing voice separation with deep U-Net convolutional networks,” in *18th International Society for Music Information Retrieval Conference*, 2017, pp. 23–27.
 - [27] C. S. J. Doire and O. Okubadejo, “Interleaved Multitask Learning for Audio Source Separation with Independent Databases,” *arXiv e-prints*, p. arXiv:1908.05182, Aug 2019.
 - [28] 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>
 - [29] I. Kavalero, S. Wisdom, H. Erdogan, B. Patton, K. Wilson, J. Le Roux, and J. R. Hershey, “Universal sound separation,” in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2019, pp. 175–179.
 - [30] Y. Luo and N. Mesgarani, “Tasnet: time-domain audio separation network for real-time, single-channel speech separation,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 696–700.
 - [31] S. Uhlich, M. Porcu, F. Giron, M. Enenkl, T. Kemp, N. Takahashi, and Y. Mitsufuji, “Improving music source separation based on deep neural networks through data augmentation and network blending,” in *Acoustics, Speech and Signal Processing (ICASSP)*, 2017 IEEE International Conference on. IEEE, 2017, pp. 261–265.
 - [32] E. M. Grais, G. Roma, A. J. Simpson, and M. Plumbley, “Combining mask estimates for single channel audio source separation using deep neural networks,” *Interspeech2016 Proceedings*, 2016.
 - [33] D. Stoller, S. Ewert, and S. Dixon, “Adversarial semi-supervised audio source separation applied to singing voice extraction,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 2391–2395.
 - [34] H. Choi, J.-h. Lee, and K. Lee, “Singing voice separation using generative adversarial networks,” in *ML4Audio Workshop, 31st Conf. Neural Information Processing Systems (NIPS 2017)*, 2017.
 - [35] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “SDR – half-baked or well done?” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 626–630.
 - [36] Z. Barzelay and Y. Y. Schechner, “Harmony in motion,” in *2007 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2007, pp. 1–8.
 - [37] Y. Aytar, C. Vondrick, and A. Torralba, “SoundNet: Learning sound representations from unlabeled video,” in *Advances in neural information processing systems*, 2016, pp. 892–900.
 - [38] R. Arandjelovic and A. Zisserman, “Look, listen and learn,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 609–617.
 - [39] —, “Objects that sound,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 435–451.
 - [40] A. Senocak, T.-H. Oh, J. Kim, M.-H. Yang, and I. S. Kweon, “Learning to localize sound sources in visual scenes: Analysis and applications,” *To appera in IEEE Transactions on Pattern Analysis and Machine Intelligence*. *arXiv preprint arXiv:1911.09649*, 2019.
 - [41] R. Gao, T.-H. Oh, K. Grauman, and L. Torresani, “Listen to look: Action recognition by previewing audio,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 10 457–10 467.
 - [42] S. Parekh, S. Essid, A. Ozerov, N. Q. Duong, P. Pérez, and G. Richard, “Weakly supervised representation learning for audio-visual scene analysis,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 416–428, 2019.
 - [43] J. Y. Liu, Y. H. Yang, and S. K. Jeng, “Weakly-Supervised Visual Instrument-Playing Action Detection in Videos,” *IEEE Transactions on Multimedia*, vol. 21, no. 4, pp. 887–901, 4 2019.
 - [44] A. Owens and A. A. Efros, “Audio-visual scene analysis with self-supervised multisensory features,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 631–648.
 - [45] A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. Wilson, A. Hassidim, W. T. Freeman, and M. Rubinstein, “Looking to listen at the cocktail party: a speaker-independent audio-visual model for speech separation,” *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, p. 112, 2018.
 - [46] R. Lu, Z. Duan, and C. Zhang, “Audio-visual deep clustering for speech separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 11, pp. 1697–1712, 2019.
 - [47] S. Parekh, A. Ozerov, S. Essid, N. Q. Duong, P. Pérez, and G. Richard, “Identify, locate and separate: Audio-visual object extraction in large video collections using weak supervision,” in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2019, pp. 268–272.
 - [48] V. Dumoulin, E. Perez, N. Schucher, F. Strub, H. d. Vries, A. Courville, and Y. Bengio, “Feature-wise transformations,” *Distill*, 2018, <https://distill.pub/2018/feature-wise-transformations>.
 - [49] O. Slizovskaia, L. Kim, G. Haro, and E. Gomez, “End-to-end sound source separation conditioned on instrument labels,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 306–310.
 - [50] E. Tzinis, S. Wisdom, J. R. Hershey, A. Jansen, and D. P. Ellis, “Improving universal sound separation using sound classification,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 96–100.

- [51] K. Schulze-Forster, C. Doire, G. Richard, and R. Badeau, "Weakly informed audio source separation," in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, 2019, pp. 268–272.
- [52] F. Pishdadian, G. Wichern, and J. L. Roux, "Finding strength in weakness: Learning to separate sounds with weak supervision," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, pp. 2386–2399, 2020.
- [53] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [54] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [55] J. F. Montesinos, O. Slizovskaia, and G. Haro, "Solos: A dataset for audio-visual music source separation and localization," in *Proceedings of the IEEE 22nd International Workshop on Multimedia Signal Processing*, 2020.
- [56] E. Vincent, R. Gribonval, and C. Fevotte, "Performance measurement in blind audio source separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 4, pp. 1462–1469, July 2006.
- [57] K. Kilgour, M. Zuluaga, D. Roblek, and M. Sharifi, "Fréchet audio distance: A reference-free metric for evaluating music enhancement algorithms," *Proc. Interspeech 2019*, pp. 2350–2354, 2019.
- [58] D. Stoller, S. Ewert, and S. Dixon, "Jointly detecting and separating singing voice: A multi-task approach," in *International Conference on Latent Variable Analysis and Signal Separation*. Springer, 2018, pp. 329–339.
- [59] A. Liutkus and F.-R. Stöter, "sigsep/norbert: First official norbert release," Jul. 2019. [Online]. Available: <https://doi.org/10.5281/zenodo.3269749>
- [60] M. Miron, J. Janer Mestres, and E. Gómez Gutiérrez, "Generating data to train convolutional neural networks for classical music source separation," in *Lokki T, Pätynen J, Välimäki V, editors. Proceedings of the 14th Sound and Music Computing Conference; 2017 Jul 5-8; Espoo, Finland. Aalto: Aalto University; 2017. p. 227-33. Aalto University, 2017.*

Olga Slizovskaia received the M.Sc. degree in computational mathematics and cybernetics from Moscow State University, Moscow, Russia, in 2013. She has been working towards the Ph.D. degree in audio-visual music information retrieval at Universitat Pompeu Fabra (UPF), Barcelona, Spain, since 2016. Her research interests include developing machine learning and deep learning methods for audio and multimodal data analysis.

Gloria Haro is an associate professor at Universitat Pompeu Fabra (UPF, Spain). She finished her Ph.D. at UPF in 2005. She was a postdoctoral fellow at the Institute for Mathematics and its Applications, University of Minnesota (USA, 2005-2007), a "Juan de la Cierva" fellow at Universitat Politècnica de Catalunya (Spain, 2007-2008), and a "Ramón y Cajal" fellow at UPF (Spain, 2008-2014). Her current research is focused on video understanding and automatic video editing tools for post-production, which involves research in basic problems such as optical flow estimation, video inpainting, segmentation, and depth estimation. She is also working in audio-visual approaches for video understanding, such as instrument recognition, source separation and sound localisation.

Emilia Gómez (Telecommunication Engineer, PhD in Computer Science) is an associate professor at Universitat Pompeu Fabra in Barcelona. She also leads the Human Behaviour and Machine Intelligence (HUMINT) project at the Centre for Advanced Studies of the Joint Research Centre, European Commission. During her career, Emilia Gómez has developed data-driven machine learning algorithms to support music listening experiences. Starting from music, she researches on the social, cultural and ethical impact of Artificial Intelligence systems and related aspects such as diversity, fairness and trust.