

---

This is an electronic reprint of the original article.  
This reprint may differ from the original in pagination and typographic detail.

Lluís, Francesc; Meyer-Kahlen, Nils; Chatziioannou, Vasileios; Hofmann, Alex  
**Direction specific ambisonics source separation with end-to-end deep learning**

*Published in:*  
Acta Acustica

*DOI:*  
[10.1051/aacus/2023020](https://doi.org/10.1051/aacus/2023020)

Published: 01/01/2023

*Document Version*  
Publisher's PDF, also known as Version of record

*Published under the following license:*  
CC BY

*Please cite the original version:*  
Lluís, F., Meyer-Kahlen, N., Chatziioannou, V., & Hofmann, A. (2023). Direction specific ambisonics source separation with end-to-end deep learning. *Acta Acustica*, 7, Article 29. <https://doi.org/10.1051/aacus/2023020>

---

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



# Direction specific ambisonics source separation with end-to-end deep learning

Francesc Lluís<sup>1,\*</sup> , Nils Meyer-Kahlen<sup>2</sup> , Vasileios Chatziioannou<sup>1</sup> , and Alex Hofmann<sup>1</sup>

<sup>1</sup>Dept. of Music Acoustics, University of Music and Performing Arts, 1030 Vienna, Austria

<sup>2</sup>Dept. for Signal Processing and Acoustics, Aalto University, Otakaari 5, 0215 Espoo, Finland

Received 13 June 2022, Accepted 2 May 2023

**Abstract** – Ambisonics is a scene-based spatial audio format that has several useful features compared to object-based formats, such as efficient whole scene rotation and versatility. However, it does not provide direct access to the individual source signals, so that these have to be separated from the mixture when required. Typically, this is done with linear spherical harmonics (SH) beamforming. In this paper, we explore deep-learning-based source separation on static Ambisonics mixtures. In contrast to most source separation approaches, which separate a fixed number of sources of specific sound types, we focus on separating arbitrary sound from specific directions. Specifically, we propose three operating modes that combine a source separation neural network with SH beamforming: refinement, implicit, and mixed mode. We show that a neural network can implicitly associate conditioning directions with the spatial information contained in the Ambisonics scene to extract specific sources. We evaluate the performance of the three proposed approaches and compare them to SH beamforming on musical mixtures generated with the musdb18 dataset, as well as with mixtures generated with the FUSS dataset for universal source separation, under both anechoic and room conditions. Results show that the proposed approaches offer improved separation performance and spatial selectivity compared to conventional SH beamforming.

**Keywords:** Source separation, Ambisonics, Deep learning, Spatial audio

## 1 Introduction

Ambisonics is a scene-based spatial audio format [1] that is widely adopted in immersive and virtual reality audio applications. It is based on representing the audio scene in the spherical harmonics (SH) domain, where each Ambisonics channel represents one SH component out of an order-limited set. This offers a universal framework for recording, transmitting and reproducing spatial audio material, with several useful features such as efficient whole scene rotation that does not depend on the number of sources in the scene, and independence between capturing and reproduction setup. This versatility comes at the cost of not having access to isolated signals of individual sound sources in the scene, as is the case in object-based spatial audio formats. Yet, access to individual sound sources or the ability to extract sound from specific directions is often required, for example, to apply modifications to the scene, or to perform analysis or further processing on individual sources. Apart from post-processing of Ambisonics music or natural scene recordings, source separation algorithms can be used for enhancing certain sound sources like human

speakers from a complete scene captured with head-worn arrays [2]. The presented algorithm would be applied after converting such a capture to the Ambisonics domain [3, 4].

The most conventional approach for extracting individual sources from an Ambisonics scene is to exploit the spatial information by applying SH beamforming. To do so, sound from a target direction is extracted by linear combination of the Ambisonics channels. Through the years, different beamformer designs have been presented [5–7], including both signal-independent and signal-dependent variants.

A completely different family of methods for extracting sound sources from a mixture is based on source separation techniques, which operate in the time-frequency domain. Source separation has been extensively studied for single-channel mixtures, but approaches have been proposed to separate sources from multichannel recordings as well, for example [8, 9]. Although using specifically Ambisonics mixtures for multichannel source separation is still relatively rare, some methods have been presented. Epain et al. [10] applied independent component analysis and Hafsati et al. [11] used local Gaussian modelling to perform source separation on Ambisonics signals. Also, several authors have proposed the use of multichannel non-negative matrix

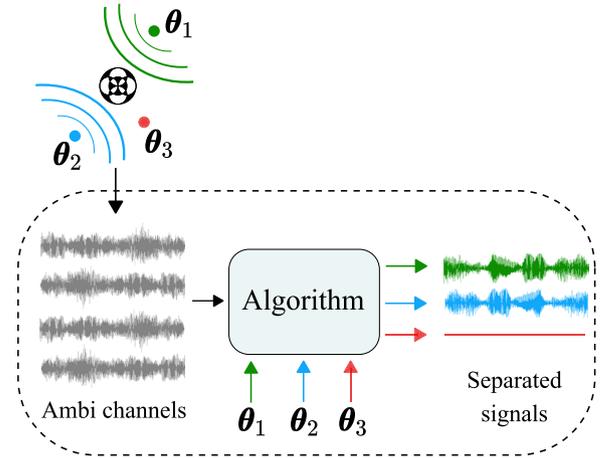
\*Corresponding author: [lluis-salvado@mdw.ac.at](mailto:lluis-salvado@mdw.ac.at)

factorization [12], and non-negative tensor factorization in several variants [13–15], operating on Ambisonics signals. In comparison to beamforming where a target direction is selected, all these source separation techniques extract sound sources of a certain type from the mixture, the number of sources needs to be known, and the complete scene is decomposed into all components. In addition, the separation stage is computationally expensive and time demanding.

Concurrently, advances in deep learning have led to large improvements in audio source separation in comparison to other methods [16, 17]. As a logical development, some approaches that combine deep learning methods and spatial processing through beamforming have already been proposed. For example, in [18, 19] a neural network is used to inform a beamformer that predicts frequency-dependent signal and noise covariance matrices to perform spatial filtering. Although encouraging results are obtained with such learning-based approaches, they are trained and tested assuming a deterministic number of sources in the mixture and closed domain sound types, such as speech signals. In contrast, our approach aims to separate any desired number of sources and type of signals. A related approach uses multichannel audio recordings and a neural network to separate speech signals in the horizontal plane [20], but it operates on microphone array data instead of Ambisonics signals and localizes sound sources itself, rather than separating sound from a specified direction.

In this paper, we adopt a data-driven approach to the task of extracting signals from specific directions given an Ambisonics mixture (as illustrated in Fig. 1). Specifically, we explore three operating modes that involve end-to-end deep learning, i.e. from waveform to waveform, and SH beamforming: (1) refinement mode, (2) implicit mode, and (3) mixed mode. In refinement mode, the deep neural network is used solely for the refinement of the single channel SH beamformer output, pointing to a target direction. In implicit mode, the Ambisonics mixture is directly provided to the network and the target direction is used to condition the network output. In mixed mode, both the Ambisonics mixture and the beamformer output are provided to the network, while the target direction is also used to condition the network output. The aim of the study is then to analyze the performance of the three modes and to compare them to conventional SH beamforming, for different Ambisonics orders. To do so, we assess the source separation performance in the direction of the sources and the capability of the methods to predict silence in regions where no source is placed, i.e. the spatial selectivity. These evaluations are performed for anechoic and room conditions, and considering musical mixtures and universal mixtures, composed by unknown number of sources from an open domain of sound types. The performance is reported considering Ambisonics mixtures with order between one and four. The code<sup>1</sup> and listening examples [21] accompanying the paper are available online.

<sup>1</sup> <https://github.com/francesclluis/direction-ambisonics-source-separation>



**Figure 1.** Illustration of an algorithm separating the sound sources from a raw Ambisonics mixture given the directions of interest. In this case, two sources (blue and green dots) located at directions  $\theta_1$  and  $\theta_2$  are emitting sound, while no sound is coming from direction  $\theta_3$ .

The paper is organized as follows. Section 2 provides background on Ambisonics and SH beamforming, which serves as a baseline in our experiments. Section 3 details the neural network architecture and the training procedure. The datasets used in the experiments are described in Section 4. Section 5 presents the evaluation metrics and the obtained results. Section 6 discusses the results and Section 7 concludes the paper.

## 2 Background

### 2.1 Ambisonics

In an ideal, instantaneous Ambisonics mixture,  $K$  far-field sound sources  $s_k(t)$  for  $k \in [1, \dots, K]$ , placed at directions  $\theta_k$  are represented as,

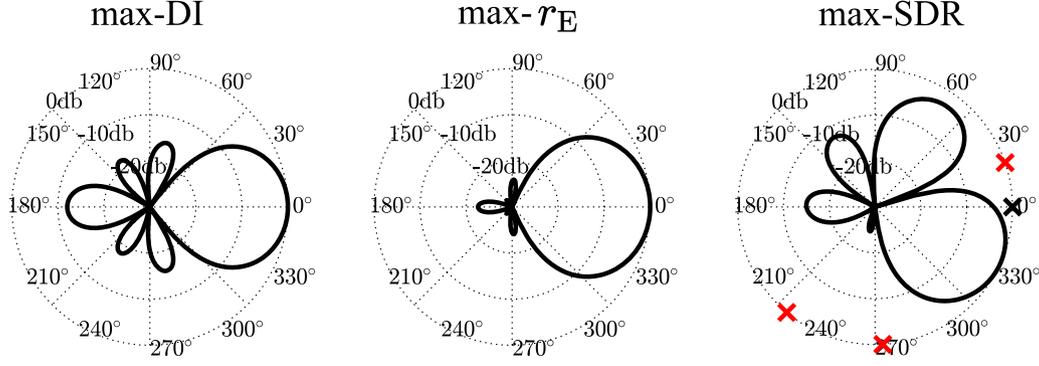
$$\mathbf{x}_N(t) = \sum_{k=1}^K s_k(t) \mathbf{y}_N(\theta_k), \quad (1)$$

where  $\mathbf{y}_N(\theta) = [Y_0^0(\theta) \ Y_1^{-1}(\theta) \ Y_1^0(\theta) \ \dots \ Y_N^N(\theta)]^\top$  is a vector of real-valued spherical harmonics, defined as,

$$Y_n^m(\theta) = N_{nm} P_n^m(\sin \vartheta) \begin{cases} \sin(|m|\phi) & m < 0 \\ 1 & m = 0, \\ \cos(m\phi) & m > 0 \end{cases} \quad (2)$$

$$N_{nm} = \sqrt{\frac{(2n+1)(n-m)!}{4\pi(n+m)!}}, \quad (3)$$

evaluated at the direction  $\theta = [\phi, \vartheta]$ , where  $0 \leq \phi < 2\pi$  is the azimuth and  $0 \leq \vartheta \leq \pi$  is the zenith angle.  $N$  is the maximal Ambisonics order and  $P_n^m$  are the associated Legendre polynomials. In acoustics, it is common to call the index  $0 \leq n \leq N$  the *order* and  $-n \leq m \leq n$  the *degree*



**Figure 2.** Signal-independent max-DI and max- $r_E$ , and signal-dependent max-SDR Beamformer pointing to a source at  $(0^\circ, 0^\circ)$  for a maximal order of  $N = 3$ . The black cross symbolizes the source and the red crosses symbolize interferers. 30 dB dynamics are shown.

of each SH component. In the full set of  $(N + 1)^2$  signals, each channel of  $\chi_N$  represents one spherical harmonic.

A convolutive mixture can be represented using SH domain directional room impulse reponses (DRIR), sometimes also called Ambisonics room impulse responses (ARIR),  $\check{h}_k(t) \in \mathbb{R}^{(N+1)^2}$ , which describe the transfer path between each source and an ideal SH receiver,

$$\chi_N(t) = \sum_{k=1}^K s_k(t) * \check{h}_k(t). \quad (4)$$

## 2.2 SH beamforming

To spatially separate sound from an Ambisonics mixture, SH beamforming uses a linear combination of the Ambisonics channels. Ideally, a beam pattern that is constant over frequency can be achieved with a real, frequency-independent weight vector  $\mathbf{d} \in \mathbb{R}^{(N+1)^2}$  such that,

$$\hat{s}_{\text{BF}}(t) = \mathbf{d}^\top \chi_N(t), \quad (5)$$

where  $\hat{s}_{\text{BF}}$  is the output of the beamformer,  $[\cdot]^\top$  is the transpose operator, and  $\chi_N$  is the Ambisonics mixture of order  $N$ . Note that the weights  $\mathbf{d}$  can be chosen arbitrarily and their choice determines the direction and shape of the implied beam-pattern.

In this work, we use two common frequency- and signal-independent SH beamformers as baselines: the beamformer with maximal directivity index (max-DI) [6, 7], which is sometimes also referred to as plane-wave decomposition, and the beamformer with maximal energy vector (max- $r_E$ ) [1]. As signal-independent beamformers, they extract sound from a specified target direction  $\theta_t$ , and the shape of the beam is derived from the global optimizatin criteria, DI, and  $r_E$  vector length. The two objectives are given by,

$$\text{DI} = 10 \log \frac{4\pi g^2(\theta_t)}{\int_{\theta} g^2(\theta) d\theta}, \quad r_E = \frac{\int_{\theta} g^2(\theta) \theta d\theta}{\int_{\theta} g^2(\theta) d\theta}, \quad (6)$$

respectively, where  $g(\theta)$  is the pattern of the beamformer evaluated at the direction  $\theta$ . It is obtained by evaluating,

$$g(\theta) = \mathbf{d}^\top \mathbf{y}_N(\theta). \quad (7)$$

Weights that optimize these global criteria are given by,

$$\mathbf{d}_{\text{max-DI}} = \mathbf{y}_N(\theta_t), \quad \mathbf{d}_{\text{max-}r_E} = \text{diag}_N(w_n) \mathbf{y}_N(\theta_t), \quad (8)$$

with the max- $r_E$  order weights, which can be approximated by ([1], p. 188),

$$w_n \approx P_n \left( \frac{\cos(137.9^\circ)}{N + 1.51} \right), \quad (9)$$

in which  $P_n$  is the Legendre function of  $n$ -th order [1], and  $\text{diag}_N(w_n)$  denotes expanding the order weights  $w_n$  to a diagonal matrix, with one weight for all the elements corresponding to one order  $n$ . See Figure 2 for the corresponding beam patterns. While the max-DI pattern has a narrow main lobe at the cost of significant side lobes, the max- $r_E$  pattern offers a compromise between main lobe width and side lobe strength.

In addition, we use another SH beamformer for comparison, which we refer to as the max-SDR beamformer. Given a particular ground truth signal of length  $T$ , now written as a vector  $\mathbf{s} = [s(0), s(1), \dots, s(T-1)]^\top$ , the max-SDR beam represents the beam pattern that extracts the signal  $\mathbf{s}$  with the maximum possible source -to-distortion (SDR) [22] ratio through SH beamforming. The SDR between a reference signal  $\mathbf{s}$  and an estimated signal  $\hat{\mathbf{s}}$  is defined as:

$$\text{SDR}(\mathbf{s}, \hat{\mathbf{s}}) = 10 \log_{10} \left( \frac{\|\mathbf{s}\|^2}{\|\mathbf{s} - \hat{\mathbf{s}}\|^2} \right). \quad (10)$$

The maximal SDR is achieved by finding the minimum squared error between ground truth and estimated sources. This is equivalent to finding the minimum mean squared error (MMSE) beamformer ([23], p. 446), given full knowledge of the source signal. The coefficients are found as,

$$\mathbf{d}_{\text{max-SDR}} = \mathbf{C}^{-1} \mathbf{X}_N^\top \mathbf{s}, \quad (11)$$

where  $\mathbf{C} = (\mathbf{X}_N^\top \mathbf{X}_N)$  is the spatial covariance matrix of the input signal of length  $T$ , which is stacked into a matrix  $\mathbf{X}_N = [\chi_N(0), \chi_N(1), \dots, \chi_N(T-1)]^\top$ .

As opposed to the other beamformers, the max-SDR pattern is signal-dependent. The approach is not directly

applicable in practical situations, as the ground truth signal  $\mathbf{s}$  is obviously not available. Also, some channels might be amplified excessively, which would introduce noise. In our investigation, the use of the max-SDR beamformer is of theoretical interest. It serves as an upper bound, to see which separation is maximum possible by frequency-independent spatial processing alone, assuming that the ground truth source signal is known. As seen in Figure 2, the max-SDR beamformer will tend to place zeros in the directions of interfering sources, if possible.

### 3 Proposed approach

We propose the use of deep learning to estimate the signal  $\mathbf{s}$  from a specific target direction  $\boldsymbol{\theta}_t$ , given the raw Ambisonics mixture  $\mathbf{X}_N$  as input signal. To this end, we explore three different operating modes which involve a deep neural network and SH beamforming.

The first is refinement mode, where the goal is to find a function  $f_1$  with the structure of a neural network such that  $f_1(\hat{\mathbf{s}}_{\text{BF}}) = \mathbf{s}$ . In this case, the neural network is used solely to refine the single channel beamformer output, which points to the target direction. Note that the neural network is not informed about the target direction and only relies on the beamformer output to enhance the beamforming operation. In addition, the beamformer output is normalized before entering the neural network as we want the proposed method to be independent of the beamformer output gain.

For the second mode, implicit mode, the goal is to find a function  $f_2$  such that  $f_2(\mathbf{X}_N, \boldsymbol{\theta}_t) = \mathbf{s}$ . In this case, the raw Ambisonics signal and the target direction are directly passed to the neural network. The training objective then forces the neural network to implicitly perform the whole beamforming operation, by learning the correspondence between the spatial information contained in the Ambisonics mixture and the conditioning direction to then perform source separation.

The third one is mixed mode, where the goal is to find a function  $f_3$  such that  $f_3(\mathbf{X}_1, \boldsymbol{\theta}_t, \hat{\mathbf{s}}_{\text{BF}}) = \mathbf{s}$ . In this case, the output of the SH beamforming is concatenated to the first order Ambisonics mixture as an extra channel. Note that the output of the SH beamforming is computed with the corresponding order but the Ambisonics mixture given to the neural network is always fixed to first order. This decision was made based on preliminary observations, similar to [24], where an increase in order of the Ambisonics input, did not substantially improve the overall performance, whereas an increase of the beamforming order did.

For all operating modes,  $\hat{\mathbf{s}}_{\text{BF}}$  corresponds to the output of a max- $\mathbf{r}_E$  beamformer at the target direction and the function  $f$  is an adapted version of the Demucs neural network architecture [25]. Specifically, Demucs input channels are modified according to the Ambisonics order and for implicit and mixed mode a global conditioning approach [20] is used to guide the separation according to the target direction. Demucs was originally designed to separate four well defined instrument types from single-channel mixture signals. Hence, the original Demucs architecture outputs as many

channels as known sources where each output corresponds to an instrument. In the present work, the output of the Demucs network is always a single channel corresponding to the audio at the target direction regardless of the number of sources in the mixture or the type of sound sources.

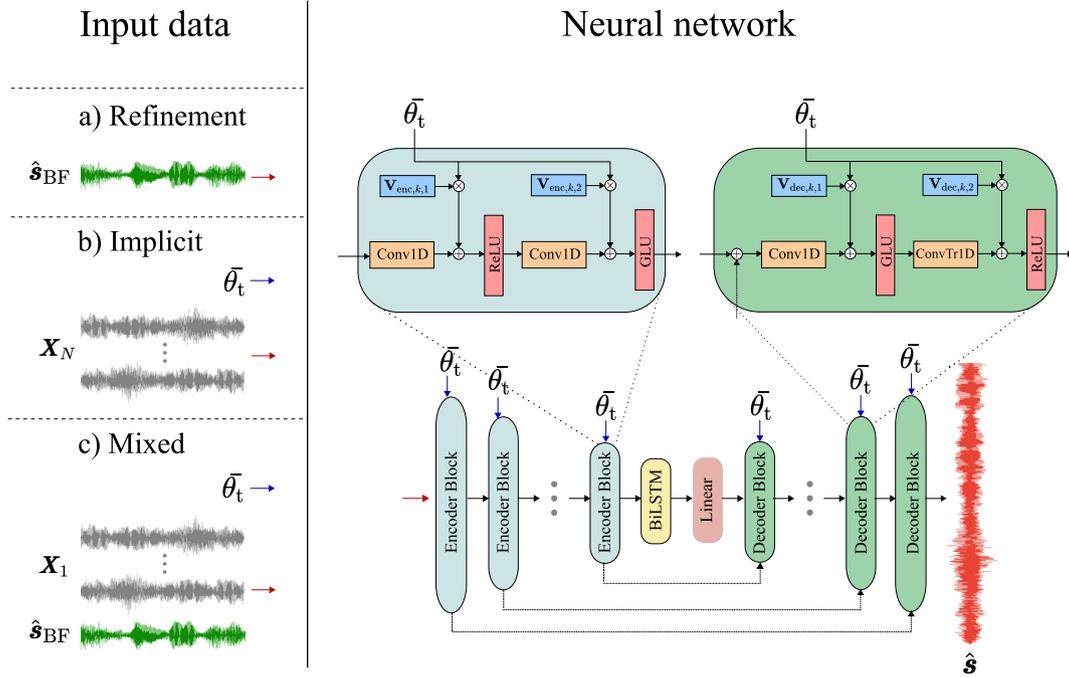
#### 3.1 Demucs architecture

Demucs [25] is a convolutional neural network that operates in the waveform domain with a U-net-like architecture [26], i.e. an encoder-decoder architecture with skip connections (see Fig. 3). The encoder-decoder structure can learn multi-resolution features of the Ambisonics mixture in the time-space domain which enables to capture the Ambisonics channel variations at different scales in both domains. The skip connections allow to propagate low level information through the network which otherwise may be lost. In this case, information related to level and phase differences between the raw Ambisonics channels can be accessed in later decoding blocks for further source separation. Each Demucs encoding block consists of an initial convolution operation that downsamples the input feature maps by applying kernels with a size of 8 and a stride of 4, while also increasing the number of channels by a factor of 2. Note that the first block is an exception as it has a fixed number of output channels set to 64. Then, a Rectified Linear Unit (ReLU) activation function is applied followed by a  $1 \times 1$  convolution with a Gated Linear Unit Activation [27] (GLU). At the bottleneck of the network, i.e. between the encoder and decoder parts, a Bidirectional Long-Short Term Memory (BiLSTM) followed by a linear operation is applied to provide long range context. The decoder part reverses the encoder process. Each decoding block first adds the skip connections from the encoder at the same level of hierarchy. Then, a  $1 \times 1$  convolution with a Gated Linear Unit Activation (GLU) is applied. Next, a transposed convolution upsamples the feature maps by applying kernels with a size of 8 and a stride of 4, while also halving the number of channels. Finally, a ReLU activation is used. Note that the last decoding block is an exception, which neither halves the number of channels nor applies the ReLU activation.

#### 3.2 Conditioning

In implicit and mixed mode, we use a global conditioning approach to inform Demucs about the target direction. Similarly to [20], the conditioning information is inserted at each block of the Demucs network after being multiplied by a learnable linear projection  $\mathbf{V}_{\cdot,q,\cdot}$ . Specifically, in this case we scale the target direction  $\boldsymbol{\theta}_t = [\phi_t, \vartheta_t]$  with azimuth angle  $\phi_t \in [-\pi, \pi]$  and zenith angle  $\vartheta_t \in [0, \pi]$ , such that the scaled target direction  $\bar{\boldsymbol{\theta}}_t = [\bar{\phi}_t, \bar{\vartheta}_t]$  is defined as  $\bar{\phi}_t \in [-1, 1]$  and  $\bar{\vartheta}_t \in [-1, 1]$ . Then the Demucs encoder and decoder takes the following expression:

$$\begin{aligned} \text{Encoder}_{q+1} &= \text{GLU}(\mathbf{W}_{\text{encoder},q,2} * \text{ReLU}(\mathbf{W}_{\text{encoder},q,1} * \\ &\quad \text{Encoder}_q + \mathbf{V}_{\text{encoder},q,1} \bar{\boldsymbol{\theta}}_t) + \mathbf{V}_{\text{encoder},q,2} \bar{\boldsymbol{\theta}}_t), \end{aligned} \quad (12)$$



**Figure 3.** Left: Input data for the following operating modes: (a) refinement, (b) implicit, (c) mixed. Right: Overview diagram of the neural network.  $\mathbf{X}_N$  refers to the  $n$ -th order Ambisonics mixture,  $\mathbf{X}_1$  is the first order Ambisonics mixture,  $\bar{\theta}_t$  is the scaled target direction,  $\hat{\mathbf{s}}_{BF}$  corresponds to the output of a max- $r_E$  beamformer at the target direction, and  $\hat{\mathbf{s}}$  is the estimated separated signal.

$$\begin{aligned} \text{Decoder}_{q-1} &= \text{ReLU}(\mathbf{W}_{\text{decoder},q,2} * \text{GLU}(\mathbf{W}_{\text{decoder},q,1} * \\ &(\text{Encoder}_q + \text{Decoder}_q) + \mathbf{V}_{\text{decoder},q,1} \bar{\theta}_t) \\ &+ \mathbf{V}_{\text{decoder},q,2} \bar{\theta}_t), \end{aligned} \quad (13)$$

where  $\text{Encoder}_{q+1}$  and  $\text{Decoder}_{q-1}$  are the outputs from the  $q$ -th level encoder and decoder blocks respectively.  $\mathbf{W}_{\cdot,q,\cdot}$  are the 1-D kernel weights at the  $q$ -th block. ReLU and GLU are the corresponding activation functions. The operator  $*$  corresponds to the 1-D convolution while  $*^\top$  denotes a transposed convolution operation, as commonly defined in the deep learning frameworks [28].

### 3.3 Supervised training

For all operating modes, the network is trained in a supervised manner. The parameters of the network are optimized to reduce the  $\ell_1$  loss between the estimated signal and the ground truth signal at the target direction. During training, as target direction, we randomly select one of the source directions and uniformly perturb it within a  $2.5^\circ$  window. This perturbation determines the spatial selectivity of the network at the inference stage. The network is trained for 200 epochs using the Adam optimizer. The learning rate is set to  $1 \times 10^{-4}$  and it is reduced by a factor 0.1 after 10 epochs with no improvement in the validation set loss. The batch size is set to 16. After the training process, we select the weights with the lowest validation loss for testing purposes. Both training and testing are conducted on a single

Titan RTX GPU. The training stage takes about 18 h while the inference takes 47 milliseconds for a single data sample (value averaged from 300 different separation predictions).

## 4 Datasets

We study the performance of the proposed methods using two different datasets: the Musdb18 [29], which contains music signals, and the Free Universal Sound Separation (FUSS) dataset [30], which contains a wide range of signals from open domain sound types. In addition, for each dataset we create an anechoic and a room version.

### 4.1 Data generation

#### 4.1.1 Musdb18

We use musical signals from the Musdb18 dataset to create training, validation, and testing data. The Musdb18 dataset contains a “train” folder with 100 songs and a “test” folder with 50 songs. For each song, the dataset provides the isolated signals of the *drums*, *bass* and *vocals* sound sources at 44.1 kHz. We use signals from 90 songs in the “train” folder to generate training data and the remaining 10 songs are used to generate validation data. For training and validation, a single example is created by first selecting six-second long audio segments from the isolated signals at a random time. Therefore, every isolated signal is taken from a random song for each source, so they do not necessarily come from the same piece. Then, to create an Ambisonics mixture, a random direction is assigned to each

of the sources, and the audio segments are encoded to up to fourth order Ambisonics and mixed using Equation (1). For the generated mixtures, it is assured that all pairs of sources are at least  $5^\circ$  great circle distance apart from each other. The great circle distance is the angle between two points on a sphere, defined as,

$$\angle(\theta_i, \theta_j) = \arctan \mathbf{x}_i^\top \mathbf{x}_j, \quad (14)$$

where  $\mathbf{x}_i = [\cos\phi_i \sin\vartheta_i, \sin\phi_i \sin\vartheta_i, \cos\vartheta_i]^\top$  is a normalized direction vector. In this work, we only consider mixtures with static sources.

Furthermore, in 30% of all created mixtures we force one source to be silent while we verify that the remaining mixtures contain active sources. The application of this pre-processing allows the data-driven approaches to learn to provide silent output when no source is present at a given direction. This is important to assure silent output when specifying directions with no active sources during inference. For training and validation we generate 10,000 and 1000 mixtures respectively. Regarding test data generation, the same encoding is applied to generate a total of 1000 mixtures using the “test” folder. In this case, single examples are created using six-second long audio segments coming from the same song at the same time and no sources are silenced.

#### 4.1.2 FUSS

The FUSS dataset was created for universal sound separation. Universal sound separation algorithms aim to separate unknown number of sources from an open domain of sound types. To this end, FUSS contains 23 hours of single-source audio data at 16 kHz drawn from 357 classes, which are used to create mixtures of one to four sources. The type of audio contained in FUSS includes natural sounds such as wind and rain, sounds of objects such as engine and alarm, and human sounds such as whistling and human voice. FUSS provides splits for training, validation, and testing with a total of 20 000, 1000, and 1000 examples respectively. We use the same partition and create six-second long Ambisonics mixtures. As previously done, for each example we assign a random direction to each of the sources, and the audio segments are encoded up to fourth order Ambisonics and mixed using Equation (1). In this case, it is assured that all pairs of sources are active and at least  $5^\circ$  great circle distance apart from each other. Note that during testing, we only consider generated mixtures that contain two or more sources.

#### 4.2 Room simulation

The anechoic Ambisonics mixtures according to (1) are a good test case, but they are far from an actual application, as usually sound sources of interest are within an enclosed room. Therefore, performance is also studied under room conditions by incorporating directional room impulse responses (DRIR) of a small room to create a convolutional mixture, as defined in (4). To create the DRIRs, sound sources are placed in a simulated room of dimensions in

the range  $(x, y, z) = (3 \pm 2 \text{ m}, 4 \pm 2 \text{ m}, 3 \pm 1 \text{ m})$ . Early reflections are simulated using the image source method with a maximal image source order of six. Wall absorption coefficients are set in octave bands, where the reflection coefficient is determined from random octave band reverberation times  $\text{RT}_f = 0.3 \pm 0.2 \text{ s}$  using Eyring’s formula [31]. For late reverberation, the response is faded over the isotropic diffuse noise at  $t_{\text{mix}} = \sqrt{V}/500 \text{ s/m}^3$ . The decay of the noise is exponentially shaped in octave bands according to the random reverberation times. Finally, although the Ambisonics mixtures contain a simulated room, we use the anechoic signals of the sound sources as ground truth for evaluating the separation performance.

## 5 Results

### 5.1 Evaluation metrics

We use two different measures of performance for evaluating the proposed methods. First, we measure the separation performance using the scale-invariant source – to – distortion ratio (SI-SDR) [32]. SI-SDR between a signal  $\mathbf{s}$  and its estimate  $\hat{\mathbf{s}}$  is defined as,

$$\text{SI-SDR}(\mathbf{s}, \hat{\mathbf{s}}) = 10 \log_{10} \left( \frac{\|\alpha \mathbf{s}\|^2}{\|\alpha \mathbf{s} - \hat{\mathbf{s}}\|^2} \right), \quad (15)$$

where  $\alpha = \text{argmin}_\alpha \|\alpha \mathbf{s} - \hat{\mathbf{s}}\|^2 = \hat{\mathbf{s}}^\top \mathbf{s} / \|\mathbf{s}\|^2$ . Specifically, for each mixture on the test set we use the ground truth direction of each active source  $\theta_k$  as the target direction for all methods. Then, the separation performance is computed between the ground truth source signal  $\mathbf{s}_k$  and the estimated signal  $\hat{\mathbf{s}}(\theta_k)$ . The SI-SDR is a metric used to evaluate the quality of separated audio sources by measuring the signal-to-noise ratio between the ground truth signal and the estimated signal, which may have an arbitrary scaling factor. An SI-SDR of 0 dB signifies that the power of the distortion is equal to the power of the ground truth signal. A positive SI-SDR value indicates that the ground truth signal has more power than the distortion, while a negative SI-SDR signifies that the power of the distortion is greater than that of the ground truth signal. Hence, higher SI-SDR values are desired.

Furthermore, we assess the performance of the models to predict silent regions in the directions where no sources are placed, i.e., its spatial selectivity. To this end, we introduce the sources-to-silence ratio (SSR). This measure is also independent of the scaling of the predicted signals. SSR is defined as,

$$\text{SSR}_k(\hat{\mathbf{s}}) = 10 \log_{10} \left( \frac{\frac{1}{K} \sum_k \|\hat{\mathbf{s}}(\theta_k)\|^2}{\frac{1}{I} \sum_t \|\hat{\mathbf{s}}(\theta_t)\|^2} \right), \quad (16)$$

where  $\{\theta_t \in \Theta \mid \angle(\theta_t, \theta_k) > 2.5^\circ\}$  are  $I$  directions from a set of 36 quasi-uniformly arranged directions on the sphere in a  $t$ -design ( $t = 8$ ),  $\Theta$  [33], excluding those that are within  $2.5^\circ$  of the target direction  $\mathbf{s}$ . Note that  $\hat{\mathbf{s}}(\theta_t)$  and  $\hat{\mathbf{s}}(\theta_k)$  are the signals predicted for the target directions  $\theta_t$  and the

ground truth source directions  $\theta_k$  respectively. Note that a SSR of 0 dB means no spatial selectivity, as it is the case for an omnidirectional receiver. A SSR of  $\infty$  dB would be achieved if silence was predicted at all directions that are at least  $2.5^\circ$  away from the sources.

## 5.2 Evaluation

We are interested in evaluating the source separation and spatial selectivity performance of all methods depending on the Ambisonics order. To this end, we compare the SI-SDR and the SSR performance considering Ambisonics mixtures with an order between one and four. In addition, we assess the performance of all methods considering the type of acoustic conditions (anechoic or room) and the type of signals in the Ambisonics mixture (music or universal). All variants are used like a traditional signal-independent beamformer, in that a target direction is provided from which sound shall be extracted, while sound from other directions shall be suppressed. Here, we provide the ground truth direction of the source to all the approaches. In the following, we report the results and behavior of the evaluated methods considering all conditions. Apart from discussing the numerical values, we also provide visualization of one example, see [Figure 4](#). Therein, the first maps show the root mean square (RMS) output of the methods, when specifying different target directions. The SI-SDR maps show the SI-SDR for the three sources present in the mix, again, when pointing to different target directions.

### 5.2.1 SH beamformer baseline

As expected, SH domain beamformers show improved performance in separation and spatial selectivity as the Ambisonics encoding-order increases, with particularly high values for high orders under anechoic conditions. For example in the case of anechoic music signals, the max- $r_E$  beamformer achieves the best separation performance for third and fourth order with a SI-SDR of 20.27 dB and 25.40 dB respectively (see [Table 1](#)). However, under room conditions, the performance of the beamformer is lower. The difference in performance of SH beamformers between anechoic and room conditions is clearly observed by the max-SDR beamformer performance in both scenarios. Very high values in the anechoic cases are due to the fact that the max-SDR approach typically places zeros in the direction of the other sources, thus cancelling them completely. Note that the max-SDR beamformer performance corresponds to the maximum possible separation by spatial processing, i.e. SH domain beamforming, alone. As discussed in detail below, the network approaches show the largest benefits over SH beamforming under the room condition. This is to be expected because deep learning based approaches can learn to remove sound reflections through non-linear operations. In contrast, the SH beamformer is strictly limited by its spatial processing resolution, only forming weighted linear combinations of the SH signals, see [Equation \(5\)](#). For low orders, linear spatial processing alone does not provide high selectivity which can be seen in the metrics, and also in

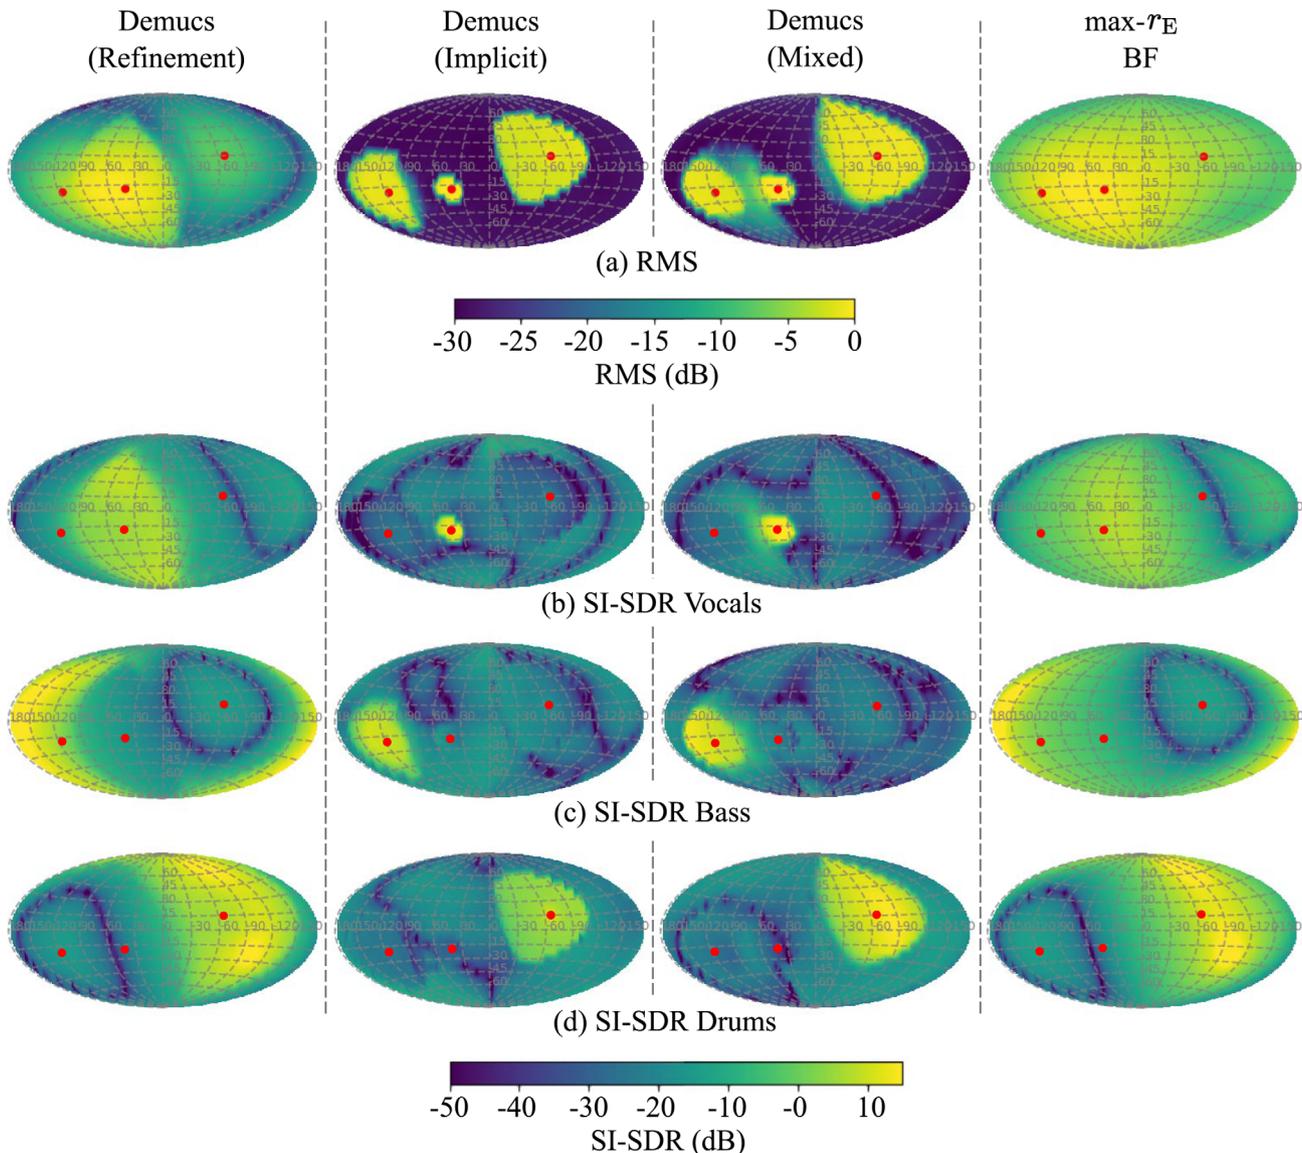
[Figure 4](#), where the output of a first order beamformer is shown in the upper right. The output power does not change strongly depending on the target direction.

### 5.2.2 Refinement mode

In refinement mode, the deep neural network is used solely for the refinement of the single channel max- $r_E$  beamformer output, pointing to the target direction. As one would expect, the refinement mode separation performance is closely related to the one achieved by the max- $r_E$  beamformer. Nevertheless, the refinement mode often improves the max- $r_E$  beamformer separation. In the cases where the max- $r_E$  beamformer separation is already high, the refinement mode achieves the best separation performance compared to other operating modes. This is the case for anechoic conditions and universal signals (see [Table 2](#)), where the refinement mode achieves a SI-SDR of 15.76 dB for third order and 21.46 dB for fourth order. However, it underperforms compared to the other methods under room conditions because the initial separation of the max- $r_E$  beamformer is not as good. Regarding spatial selectivity, the refinement mode fails to predict silence in source-free regions. The neural network cannot determine if the input sound coming from the max- $r_E$  beamformer output should be silenced or enhanced without the target direction information. The bad spatial selectivity performance is especially notable in anechoic conditions where, for example, it achieves a SSR of  $-0.28$  dB for fourth order when using music signals.

### 5.2.3 Implicit mode

In the anechoic case, the implicit mode achieves similar separation and spatial selectivity independent of the encoding order it has been trained and tested on. For music signals, it achieves a SI-SDR of 15.66 dB and a SSR of 10.70 dB for first order mixtures while it achieves a SI-SDR of 16.16 dB and a SSR of 10.08 dB for fourth order mixtures. Hence, low orders are enough for the implicit mode to learn the correspondence between the spatial information contained in the Ambisonics signal and the conditioning direction, and higher order input channels do not have large benefits. The same behaviour is observed when the implicit mode is trained using universal signals. When the implicit mode is trained and tested under room conditions, results show improved performance in separation and spatial selectivity as the encoding order is increased. For instance with music signals, it achieves a SI-SDR of 0.57 dB and a SSR of 4.90 dB for first order mixtures as opposed to 6.32 dB and 7.43 dB achieved for fourth order mixtures (see [Table 1](#)). Overall, the implicit mode is the operating mode that achieves the best SSR results out of all other operating modes for every Ambisonics order and type of signals. This behaviour can also be seen in [Figure 4](#). The network manages to strongly suppress the signal in directions that are not close to sources. In addition, it achieves competitive separation performance, especially for music signals under room conditions (see [Table 1](#)).



**Figure 4.** Visualization of (a) RMS and (b), (c), (d) SI-SDR for audio predicted in a discrete set of equiangularly distributed directions ( $100^\circ \times 50^\circ$ ) by several methods. The predictions are made using a first order Ambisonics mixture from the Musdb18 test set in anechoic conditions. The red dots correspond to the location of the sources. Corresponding listening examples can be found online [21].

#### 5.2.4 Mixed mode

In the mixed mode, both the first order Ambisonics signal and the beamformer output are provided to the network, while the target direction is also used to condition the network output. Overall, the mixed mode achieves competitive SI-SDR results compared to all other operating modes. For the more realistic case, i.e. with universal signals and room conditions, the mixed mode achieves the best separation performance compared to all other methods for any Ambisonics order (see Table 2). For first order musical mixtures in the room, the mixed mode achieves a SI-SDR of 0.76 dB while the max-SDR beamformer achieves a SI-SDR of 0.58 dB. This means that the mixed mode separation performance is similar to the maximum possible

achieved by spatial processing alone. Regarding spatial selectivity under room conditions, the mixed mode offers a similar performance independently of the encoding order. For example using universal signals, it achieves a SSR of 4.65 dB for first order and 4.91 dB for fourth order. For anechoic conditions, the SSR values show an unexpected behaviour of the model, which performs competitively for lower orders while it fails for higher orders. This can be seen using universal signals in Table 2, where it achieves a SSR of 6.61 dB for the first order and a SSR of  $-0.13$  dB for the fourth order. We suspect that this network behaviour is caused by not learning the correspondence between the target angle and the mixture, and just basing the predictions on the provided max- $r_E$  beamformer output. This leads to good separation only in anechoic and high-order scenarios, where

**Table 1.** SI-SDR and SSR median scores in dB, along with their 95% confidence interval within braces, calculated for the test set using music signals for both anechoic and room conditions. The highest performances are highlighted using bold font.

	Ambi. order	Demucs (refinement)		Demucs (implicit)		Demucs (mixed)		max-DI BF		max- $r_E$ BF		max-SDR BF
		SI-SDR	SSR	SI-SDR	SSR	SI-SDR	SSR	SI-SDR	SSR	SI-SDR	SSR	SI-SDR
Musdb18 Anechoic	1	7.50	-0.30	15.66	<b>10.70</b>	<b>16.45</b>	8.79	4.48	2.71	4.40	2.48	53.83
		(0.47)	(0.08)	(0.25)	<b>(0.21)</b>	<b>(0.28)</b>	(0.18)	(0.22)	(0.08)	(0.22)	(0.08)	(0.20)
	2	14.89	0.01	15.44	<b>10.91</b>	<b>20.62</b>	8.13	10.69	5.09	12.20	4.69	57.73
		(0.51)	(0.07)	(0.24)	<b>(0.17)</b>	<b>(0.40)</b>	(0.15)	(0.22)	(0.06)	(0.23)	(0.08)	(0.19)
	3	20.16	0.16	11.85	<b>14.94</b>	18.95	0.94	14.76	7.29	<b>20.27</b>	6.52	59.33
		(0.54)	(0.08)	(0.12)	<b>(0.27)</b>	(0.47)	(0.04)	(0.22)	(0.04)	<b>(0.22)</b>	(0.06)	(0.20)
	4	24.39	-0.28	16.16	<b>10.08</b>	23.00	-0.12	17.80	9.17	<b>25.40</b>	8.31	60.29
		(0.54)	(0.04)	(0.26)	<b>(0.13)</b>	(0.49)	(0.04)	(0.22)	(0.03)	<b>(0.23)</b>	(0.04)	(0.21)
Musdb18 Room	1	-1.16	0.20	0.57	<b>4.90</b>	<b>0.76</b>	3.85	-1.36	2.03	-1.33	1.76	0.58
		(0.33)	(0.08)	(0.25)	<b>(0.15)</b>	<b>(0.26)</b>	(0.16)	(0.22)	(0.07)	(0.22)	(0.07)	(0.20)
	2	2.94	1.14	<b>3.57</b>	<b>6.14</b>	2.97	0.86	1.83	3.73	1.50	3.38	4.95
		(0.38)	(0.09)	<b>(0.19)</b>	<b>(0.13)</b>	(0.34)	(0.10)	(0.22)	(0.07)	(0.23)	(0.08)	(0.19)
	3	4.79	1.81	<b>5.31</b>	<b>6.88</b>	5.03	1.98	4.04	5.23	3.51	4.77	8.39
		(0.36)	(0.10)	<b>(0.19)</b>	<b>(0.11)</b>	(0.25)	(0.12)	(0.22)	(0.06)	(0.23)	(0.07)	(0.20)
	4	3.87	1.89	6.32	<b>7.43</b>	<b>6.43</b>	1.43	5.53	6.45	5.07	5.93	11.05
		(0.37)	(0.09)	(0.18)	<b>(0.13)</b>	<b>(0.21)</b>	(0.11)	(0.22)	(0.07)	(0.23)	(0.07)	(0.21)

**Table 2.** SI-SDR and SSR median scores in dB, along with their 95% confidence interval within braces, calculated for the free universal sound separation test set (2–4 sources) for both anechoic and room conditions. The highest performances are highlighted using bold font.

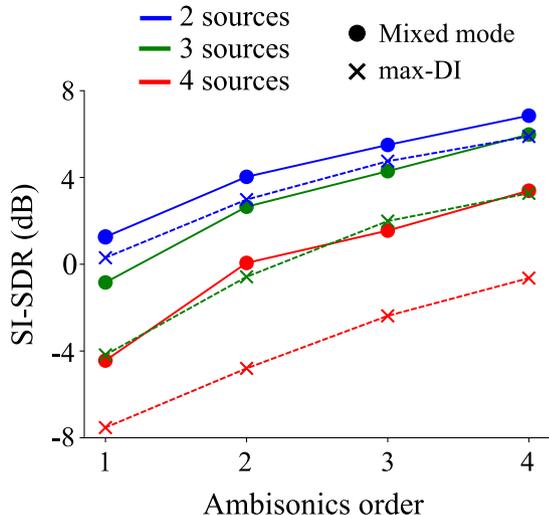
	Ambi. order	Demucs (refinement)		Demucs (implicit)		Demucs (mixed)		max-DI BF		max- $r_E$ B		max-SDR BF
		SI-SDR	SSR	SI-SDR	SSR	SI-SDR	SSR	SI-SDR	SSR	SI-SDR	SSR	SI-SDR
FUSS Anechoic	1	2.73	-0.01	9.93	<b>8.26</b>	<b>10.20</b>	6.61	2.47	3.39	2.61	2.99	40.94
		(0.70)	(0.14)	(0.48)	<b>(0.26)</b>	<b>(0.54)</b>	(0.27)	(0.66)	(0.19)	(0.67)	(0.20)	(0.43)
	2	10.99	-0.01	12.27	<b>7.89</b>	<b>15.88</b>	7.80	7.49	6.53	9.58	5.95	42.44
		(0.83)	(0.15)	(0.47)	<b>(0.27)</b>	<b>(0.58)</b>	(0.28)	(0.64)	(0.22)	(0.73)	(0.23)	(0.44)
	3	<b>15.76</b>	-0.08	13.15	7.86	15.37	-0.06	11.27	<b>8.83</b>	15.28	8.19	43.39
		<b>(0.82)</b>	(0.14)	(0.45)	(0.26)	(0.77)	(0.14)	(0.67)	<b>(0.21)</b>	(0.81)	(0.22)	(0.45)
	4	<b>21.46</b>	-0.10	12.98	8.94	20.22	-0.13	14.91	<b>10.70</b>	20.25	10.01	44.17
		<b>(0.76)</b>	(0.16)	(0.43)	(0.25)	(0.70)	(0.15)	(0.66)	<b>(0.23)</b>	(0.78)	(0.22)	(0.45)
FUSS Room	1	-4.22	2.01	-1.50	<b>5.08</b>	<b>-1.27</b>	4.65	-4.08	2.36	-4.17	2.08	-1.75
		(0.60)	(0.19)	(0.44)	<b>(0.23)</b>	<b>(0.48)</b>	(0.23)	(0.46)	(0.17)	(0.46)	(0.16)	(0.41)
	2	-0.98	2.80	1.62	<b>7.16</b>	<b>2.22</b>	3.22	-0.91	4.25	-1.22	3.95	2.88
		(0.63)	(0.22)	(0.36)	<b>(0.20)</b>	<b>(0.39)</b>	(0.21)	(0.44)	(0.19)	(0.44)	(0.19)	(0.35)
	3	1.54	3.72	3.00	<b>7.47</b>	<b>3.75</b>	4.01	1.59	5.76	0.99	5.37	6.40
		(0.53)	(0.23)	(0.33)	<b>(0.22)</b>	<b>(0.37)</b>	(0.19)	(0.43)	(0.20)	(0.42)	(0.19)	(0.34)
	4	3.61	4.43	3.59	<b>7.62</b>	<b>5.47</b>	4.91	2.84	6.77	2.65	6.46	9.04
		(0.50)	(0.24)	(0.37)	<b>(0.23)</b>	<b>(0.34)</b>	(0.21)	(0.41)	(0.20)	(0.40)	(0.20)	(0.34)

the max- $r_E$  beamformer output already provides a meaningful solution according to the training loss function.

### 5.2.5 Number of sources in the mixture

We are also interested in studying the source separation performance depending on the number of sources in the mixture. Based on the results in Table 2, we report the performance considering the best learning – based and SH beamformer methods for universal signals under room conditions, i.e. the mixed mode and the max-DI beamformer.

Figure 5 shows that for a given encoding order, the difference between the SI-SDR achieved by the mixed mode compared to the max-DI beamformer increases with the number of sources in the mixture. For a given number of sources, the improvement of SI-SDR, between the mixed mode and the max-DI beamformer, is more or less constant independently of the encoding order. This analysis indicates that as the number of sources within the mixture increases, the utilization of deep learning for source separation becomes more beneficial in comparison to the beamformer approach. However, it should be noted that the increment



**Figure 5.** FUSS room SI-SDR median results reported for the different number of sources (2, 3, 4) in the Ambisonics mixture for both mixed mode and max-DI beamformer target direction.

in performance remains consistent across the Ambisonics encoding orders evaluated.

## 6 Discussion

The results in Section 5 indicate that each of the presented operating modes has its own benefits, depending on the application scenario.

The refinement mode is adequate when the separation achieved by the SH beamformer, which is provided as input to the network, is already high. For low orders, the refinement mode performs poorly because the SH beamformer does not strongly cancel interfering sources. In this case, the network can not determine which source to refine, as it does not have access to any information concerning which part of the signal is the target, besides level differences. The low SSR values achieved by the refinement mode show that the network is not able to predict silence. It will refine any source at the output of the SH beamformer, even when it is pointing into a source-free region.

The results achieved by the implicit mode show that the neural network can implicitly learn the correspondence between the spatial information contained in the Ambisonics mixture and the target direction. It correctly interprets the target direction and extracts the sound from that direction. This also becomes evident in the SI-SDR maps shown in Figure 4, where, for the same input mixture, the neural network predictions are different for each of the conditioning target directions and correlate better with the ground truth signals in the vicinity of the source location.

Under room conditions, the implicit mode also performs better separation than common SH beamformers. This is also true for low orders in anechoic conditions. Interestingly, the implicit mode does not benefit from higher orders in anechoic mixtures. The largest advantage of the implicit mode over conventional SH beamforming lies in the higher

spatial selectivity. Pointing to a target direction where no source is present yields very low outputs. This can be seen in the RMS map shown in Figure 4. Note that this capability could potentially be used for source localization as well, in an algorithm more similar to [20], where separation and localization are combined. It also means that the separation stage is more sensitive to direction of arrival mismatches than the SH approaches, i.e., if a sound source is too far away from the provided target direction, the network would output silence. Note that the angular range in which the network outputs signal can be controlled in the training stage, by changing the perturbation applied to the target direction and the minimal distance between sources.

Overall, this shows that a source separation network conditioned only on the target direction can perform the whole beamforming operation, i.e. extract sound from specific directions, with mixtures containing an unknown number of sources and arbitrary sound type.

If the goal is to perform source separation where spatial selectivity is less important, the mixed mode might be the most useful. It provides best separation performance for first and second order in the anechoic music dataset and for all orders in the the FUSS dataset under room conditions. Testing the mixed mode with the FUSS dataset under room conditions shows the largest improvements over SH beamforming for relatively complex scenes with four sources. The increase in SI-SDR over the implicit mode comes at the cost of lower SSR, which can also be seen in Figure 4, where the RMS at directions different from the sources is higher.

Interestingly, neither the implicit mode, nor the mixed mode, which has the max- $r_E$  beamformer as input, achieves higher SI-SDR in the anechoic cases than the conventional max- $r_E$ . We see this mainly as evidence for the fact that SH beamforming under anechoic conditions is very effective, and any impairment caused by the network will be enough to result in a lower SI-SDR. The high separation performance of high order conventional SH beamforming is quickly lost in the more realistic case, when room reflections are present. For the first order room condition, the mixed mode achieves higher separation than SH beamformers. It is in a similar range as the max-SDR beamformer, which represents the maximal SDR achievable by frequency-independent SH beamforming given knowledge of the ground truth signal.

In the future, the evaluation of deep learning based methods may be extended to recordings with real microphone arrays. There, the maximal Ambisonics order is not obtained at low frequencies, due to necessary regularization [1]. This might have more impact on the performance of conventional approaches compared to deep learning ones, given that the performance of conventional SH beamforming strongly depends on the order. However, it should be tested if a model trained with the data from one microphone array can generalize to data from a different array.

Furthermore, when comparing the objective metrics for the three modes against the SH beamformer, it is important to note that higher separation may come at the cost of time-frequency artifacts, that can occur in the neural network

output, but not in SH beamformers. Ideally, a formal listening experiment should be conducted in the future. For now, listening examples are provided online that can give an impression of the perceptual quality [21]. Although the results are promising, the above artifacts could become audible in a post-processing application. Nevertheless, increasing the level of a particular signal might be feasible, without significantly reducing its quality.

## 7 Conclusions

In this paper we proposed the use of end-to-end deep learning as an alternative to conventional SH beamforming. We have shown and analyzed three different operation modes: (1) refinement, (2) implicit and (3) mixed. Specifically, the implicit and mixed mode show that a source separation network can learn associations between a target direction and the information contained in an Ambisonics scene. This allows for using such a network as one would use a beamformer, specifying a target direction and separating arbitrary sounds without adapting the training to a specific number or type of source.

The results show that, under anechoic conditions, the largest separation improvement of the proposed approaches with respect to SH beamforming is achieved for lower Ambisonics orders. In addition, better spatial selectivity is provided for all orders. Under room conditions, the application of deep learning increases both separation and spatial selectivity for all orders. Generally, the behaviour of each operating mode is similar when trained and tested using musical signals or arbitrary signals from a dataset for universal source separation.

## Data availability statement

Listening examples are available online at: <http://research.spa.aalto.fi/publications/papers/acta22-sss/> [21]. The code accompanying the paper is provided at: <https://github.com/francesclluis/direction-ambisonics-source-separation>. The dataset of Musdb18, which contains signals are available online at <https://doi.org/10.5281/zenodo.3338373> [29].

## Conflict of interest

The authors declare that they have no conflicts of interest in relation to this article.

## Acknowledgments

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 812719.

## References

1. F. Zotter, M. Frank: Ambisonics: A practical 3D audio theory for recording, studio production, sound reinforcement, and virtual reality, ser. Springer Topics in Signal Processing, Vol. 19. Springer International, 2019.
2. P. Guiraud, S. Hafezi, P.A. Naylor, A.H. Moore, J. Donley, V. Tourbabin, T. Lunner: An introduction to the speech enhancement for augmented reality (spear) challenge, in 2022 International Workshop on Acoustic Signal Enhancement (IWAENC), 5–8 September 2022, Bamberg, Germany, 2022.
3. J. Ahrens, H. Helmholtz, D.L. Alon, S.V.A. Gari: Spherical harmonics decomposition of a sound field based on microphones around the circumference of a human head, in IEEE Workshop on Applications of Signal Processing to Audio and Acoustics, 17–20 October 2021, New Paltz, NY, USA, 2021.
4. L. McCormack, A. Politis, R. Gonzalez, T. Lokki, V. Pulkki: Parametric ambisonic encoding of arbitrary microphone arrays. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 30 (2022) 2062–2075.
5. H. Teutsch: Modal array signal processing: principles and applications of acoustic wavefield decomposition, Vol. 348. Springer, 2007.
6. B. Rafaely: Fundamentals of spherical array processing. Springer Berlin Heidelberg, New York, NY, 2014.
7. D.P. Jarrett, E.A. Habets, P.A. Naylor: Theory and applications of spherical microphone array processing, ser. Springer Topics in Signal Processing, Vol. 9. Springer International Publishing, Cham, 2017. [Online]. Available: <http://link.springer.com/10.1007/978-3-319-42211-4>.
8. A.A. Nugraha, A. Liutkus, E. Vincent: Multichannel audio source separation with deep neural networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 24, 9 (2016) 1652–1664.
9. A. Ozerov, C. Fevotte: Multichannel nonnegative matrix factorization in convolutive mixtures for audio source separation. *IEEE Transactions on Audio, Speech, and Language Processing* 18, 3 (2009) 550–563.
10. N. Epain, C.T. Jin: Independent component analysis using spherical microphone arrays. *Acta Acustica United with Acustica* 98, 1 (2012) 91–102.
11. M. Hafsati, N. Epain, R. Gribonval, N. Bertin: Sound source separation in the higher order ambisonics domain, in DAFx 2019 – 22nd International Conference on Digital Audio Effects, September 2019, Birmingham, United Kingdom, 2019, pp. 1–7.
12. J. Nikunen, A. Politis: Multichannel NMF for source separation with ambisonic signals, in 2018 16th International Workshop on Acoustic Signal Enhancement (IWAENC), 17–20 September 2018, Tokyo, Japan. IEEE, 2018, p. 251255.
13. A.J. Munoz-Montoro, J.J. Carabias-Orti, P. Vera-Candeas: Ambisonics domain singing voice separation combining deep neural network and direction aware multichannel NMF, in 2021 IEEE 23rd International Workshop on Multimedia Signal Processing (MMSP), 6–8 October 2021, Tampere, Finland. IEEE, 2021.
14. Y. Mitsufuji, N. Takamune, S. Koyama, H. Saruwatari: Multichannel blind source separation based on evanescent-region-aware non-negative tensor factorization in spherical harmonic domain. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 29 (2021) 607–617.
15. M. Guzik, K. Kowalczyk: Wishart localization prior on spatial covariance matrix in ambisonic source separation using non-negative tensor factorization, in ICASSP 2022–2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 23–27 May 2022, Singapore. IEEE, 2022, pp. 446–450.

16. Y. Mitsufuji, G. Fabbro, S. Uhlich, F.-R. Stoter: Music demixing challenge 2021. *Frontiers in Signal Processing 1* (2022) 18.
17. M. Cobos, J. Ahrens, K. Kowalczyk, A. Politis: An overview of machine learning and other data-based methods for spatial audio capture, processing, and reproduction. *EURASIP Journal on Audio, Speech, and Music Processing 2022*, 1 (Dec. 2022) 1–21. [Online]. Available: <https://asmp-urasip-journals.springeropen.com/articles/10.1186/s13636-022-00242-x>.
18. A. Bosca, A. Guerin, L. Perotin, S. Kitic: Dilated U-net based approach for multichannel speech enhancement from first-order ambisonics recordings, in 2020 28th European Signal Processing Conference (EUSIPCO), 18–21 January 2021, Amsterdam, Netherlands. IEEE, 2020, pp. 216–220.
19. T. Ochiai, M. Delcroix, R. Ikeshita, K. Kinoshita, T. Nakatani, S. Araki: Beam-TasNet: Time-domain audio separation network meets frequency-domain beamformer, in ICASSP 2020–2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 4–8 May 2020, Barcelona, Spain. IEEE, 2020, pp. 6384–6388. [Online]. Available: <https://ieeexplore.ieee.org/document/9053575/>.
20. T. Jenrunggrot, V. Jayaram, S. Seitz, I. Kemelmacher-Shlizerman: The cone of silence: Speech separation by localization. *Advances in Neural Information Processing Systems 33* (2020) 20,925–20,938.
21. Online listening examples, <http://research.spa.aalto.fi/publications/papers/acta22-sss/>.
22. E. Vincent, H. Sawada, P. Bofill, S. Makino, J.P. Rosca: First stereo audio source separation evaluation campaign: data, algorithms and results, in *International Conference on Independent Component Analysis and Signal Separation*, 9–12 September 2007, London, United Kingdom. Springer, 2007, pp. 552–559.
23. H.L. Van Trees, *Detection, estimation, and modulation theory. 4: Optimum array processing*. Wiley, New York, NY, 2002.
24. F. Lluís, N. Meyer-Kahlen, V. Chatziioannou, A. Hofmann: A deep learning approach for angle specific source separation from raw ambisonics signals, in DAGA, 21–24 March 2022, Stuttgart, Germany. 2022.
25. A. Défossez, N. Usunier, L. Bottou, F. Bach: Music source separation in the waveform domain. 2019, ArXiv preprint: [arXiv:1911.13254](https://arxiv.org/abs/1911.13254).
26. O. Ronneberger, P. Fischer, T. Brox: U-net: Convolutional networks for biomedical image segmentation, in *International Conference on Medical image computing and computer-assisted intervention*, October 5–9, 2015, Munich, Germany. Springer, 2015, pp. 234–241.
27. Y.N. Dauphin, A. Fan, M. Auli, D. Grangier: Language modeling with gated convolutional networks, in *International Conference on Machine Learning*, PMLR, August 2017, Sydney, Australia. 2017, pp. 933–941.
28. A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala: Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems 32* (2019) 8026–8037.
29. Z. Rafii, A. Liutkus, F.-R. Stöter, S.I. Mimitakis, R. Bittner: Musdb18-hq – an uncompressed version of musdb18. Aug. 2019. [Online]. Available: <https://doi.org/10.5281/zenodo.3338373>.
30. S. Wisdom, H. Erdogan, D.P. Ellis, R. Serizel, N. Turpault, E. Fonseca, J. Salamon, P. Seetharaman, J.R. Hershey: What’s all the fuss about free universal sound separation data?, in ICASSP 2021–2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Jun 2021, Toronto/Virtual, Canada. IEEE, 2021, pp. 186–190.
31. H. Kuttruff: *Room acoustics*, 6th ed. CRC Press, Boca Raton, NY, 2017.
32. J. Le Roux, S. Wisdom, H. Erdogan, J.R. Hershey: SDR – half-baked or well done?, in 2019 ICASSP, 12–17 May 2019, Brighton, United Kingdom. IEEE, 2019, pp. 626–630.
33. R.H. Hardin, N.J.A. Sloane: McLaren’s improved snub cube and other new spherical designs in three dimensions. *Discrete and Computational Geometry 15* (1996) 429–441.

**Cite this article as:** Lluís F. Meyer-Kahlen N. Chatziioannou V. & Hofmann A. 2023. Direction specific ambisonics source separation with end-to-end deep learning. *Acta Acustica*, 7, 29.